

# Sound-Based Recognition of Touch Gestures and Emotions for Enhanced Human-Robot Interaction

Yuanbo Hou\*, Qiaoqiao Ren<sup>†</sup>, Wenwu Wang<sup>‡</sup>, Dick Botteldooren\*

\*WAVES, Ghent University, Belgium <sup>†</sup>AIRO, Ghent University-IMEC, Belgium <sup>‡</sup>CVSSP, University of Surrey, UK

**Abstract**—Emotion recognition and touch gesture decoding are crucial for advancing human-robot interaction (HRI), especially in social environments where emotional cues and tactile perception play important roles. However, many humanoid robots, such as Pepper, Nao, and Furhat, lack full-body tactile skin, limiting their ability to engage in touch-based emotional and gesture interactions. In addition, vision-based emotion recognition methods usually face strict GDPR compliance challenges due to the need to collect personal facial data. To address these limitations and avoid privacy issues, this paper studies the potential of using the sounds produced by touching during HRI to recognise tactile gestures and classify emotions along the arousal and valence dimensions. Using a dataset of tactile gestures and emotional interactions from 28 participants with the humanoid robot Pepper, we design an audio-only lightweight touch gesture and emotion recognition model with only 0.24M parameters, 0.94MB model size, and 0.7G FLOPs. Experimental results show that the proposed sound-based touch gesture and emotion recognition model effectively recognises the arousal and valence states of different emotions, as well as various tactile gestures, when the input audio length varies. The proposed model is low-latency and achieves similar results as well-known pretrained audio neural networks (PANNs), but with much smaller FLOPs, parameters, and model size.

**Index Terms**—Affective computing, emotion classification, touch gestures, humanoid robots, multi-temporal resolution CNN.

## I. INTRODUCTION

Robots’ perception of tactile gestures and emotions is integral to the development of advanced human-robot interaction (HRI), especially in social environments, where understanding and responding to emotional cues is critical for meaningful interactions [1]–[3]. Moreover, real-time emotion states and touch gesture recognition can significantly improve the naturalness and effectiveness of HRI [4] [5].

Previous work [6] has attempted to decode touch gestures from tactile signals, but a major limitation is that many robots lack skin that can sense touch throughout the body. Most research on humanoid robotics values touch less than vision or audio signals. Robots like Pepper [7], NAO [8], and Furhat [9] do not have extensive tactile sensors covering the entire body [10]. This limitation hinders the development of HRI systems that decode gestures and emotions based on tactile perception. To this end, tactile sensors embedded in the surface of robots are used to decode touch gestures based on pressure patterns, thus providing a direct tactile-based interaction method [11]. However, tactile sensors [12] can be intrusive and are often limited by the need for different body contacts.

In addition to tactile-based gesture recognition, most gesture decoding studies adopt vision-based methods, which use depth

cameras [13] or optical sensors [14] to capture human gestures and emotions. These vision-based models perform well in controlled environments. However, they often face significant challenges in dynamic real-world environments [15], e.g., occlusion of certain parts of the face/body or changes in lighting conditions can seriously affect the effectiveness of the vision-based models [16]. Moreover, the vision-based models usually require relatively high computational resources, and collection of training data. However, it is difficult to avoid issues related to privacy and general data protection regulation (GDPR) during data collection. These limitations restrict the scalability and practicality of vision-based methods.

To recognize gestures and emotions, researchers have also explored auditory signals and demonstrated the feasibility of sensing touch gestures on the surface of mobile devices based on acoustic signals [17]. Auditory stimuli are viewed to be potent triggers of affective responses, and specific sounds, e.g., tapping sounds, can affect people’s emotional states and behaviours [18]. Moreover, speech-based emotion recognitions have been successfully applied to adjust the behaviour of social and educational robots to adapt to current social emotions. Furthermore, the recognition of emotions such as happiness, sadness, and anger has been widely carried out using features such as the rhythm, pitch, and intensity of speech [19]–[21].

Emotions are typically analysed along two independent dimensions (arousal and valence), like in Russell’s model [22] [23], and neuroimaging studies [24] [25] support these two-dimensional representations. Hence, the arousal-valence dimensional model (AVDM) is more commonly used for sound-related affective computing, e.g., soundscape studies [26] [27], than the discrete emotion model for single discrete entities [28] [29]. Hence, this paper uses AVDM to analyse touch-related emotions.

This paper makes the first attempt to recognise touch gestures and emotions based on sounds produced during HRI. Although previous studies have shown the feasibility of sensing touch gestures based on acoustic signals, they did not further study the decoding of different gestures from participants. Whereas most of the previous auditory-based human emotion recognition related to robots is based on speech signals, this paper utilises the non-speech sounds produced during HRI.

## II. BACKGROUND

Pepper is an interactive robot developed by SoftBank Robotics [7], as shown in Fig. 1; more technical specifications,




Fig. 1: The robot Pepper’s physical information<sup>1</sup>.

please see here<sup>1</sup>. In the data collection experiment, participants are asked to express 6 gestures (*hold, pat, poke, tickle, tap, rub*) and 10 emotions by touching Pepper’s left forearm with spontaneous movements. The sounds produced by these touch movements are recorded by a microphone (the black device in Fig. 2) to form the sound dataset in this paper.

According to the Circumplex Model [23], the distribution of the 10 emotions involved in arousal-valence dimensions is shown in Fig. 3. *Happiness* and *surprise* occupy the high arousal, positive valence quadrant (Q1); *anger, fear* and *disgust* are located in the high arousal, negative valence quadrant (Q2); *sadness* and *confusion* are located in the low arousal, negative valence quadrant (Q3); *comfort* and *calming* belong to the low arousal, positive valence quadrant (Q4); the neutral emotion, *attention*, is located at the origin (Q0).




Fig. 2: The participant interacts with the robot Pepper.

### III. PROPOSED METHOD

Our ultimate goal is to develop a sound-based model that can be embedded into a robot to perceive touch gestures and emotions. To this end, this section first analyzes the limited availability of computing resources in the robot brain system and then designs a lightweight model tailored for the robot.

#### A. Robot’s brain system computing resources

Pepper’s Brain system consists of an Intel ATOM® E3845 processor and 4 GB of DDR3 RAM. The documentation<sup>2</sup> provided by Intel shows that E3845 processor’s Floating Point of Operations (FLOPs) is 11.46G per second. To deploy a model in Pepper’s brain system, the model’s FLOPs should not be greater than 11.46G; otherwise, it cannot be run on Pepper. Considering this restriction, the model to be deployed should meet the following requirements: 1) lightweight with few parameters; 2) small FLOPs and low processing latency; 3) able to process varying-length audio clips. Next, we will use these three points as guidelines to design the model.

<sup>1</sup><https://support.aldebaran.com/support/solutions/articles/80000958735-pepper-technical-specifications>

<sup>2</sup><https://www.intel.com/content/dam/support/us/en/documents/processors/APP-for-Intel-Atom-Processors.pdf>

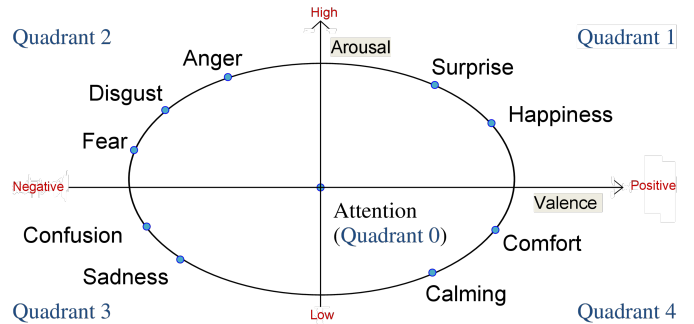


Fig. 3: Circumplex Model [23] with 10 emotions in this paper.

#### B. The proposed lightweight model: MTRCNN

Sounds caused by various touch gestures and emotions may have different durations, e.g., the sound of an angry hit and the rust of a calm touch. Thus, we propose a multi-temporal resolution convolutional neural network (MTRCNN) in Fig. 4. The convolution (Conv) part of MTRCNN consists of 3 branches with Conv kernel sizes of (3, 3), (5, 5), and (7, 7), respectively. Different kernel sizes extract representations with different resolutions. To obtain a larger Conv receptive field size (RFS) with fewer parameters, the dilated Conv [30] is used. Moreover, to avoid the gridding artifacts [31] of the dilated Conv, hybrid dilated Conv scheme [30] is adopted, so the number of filters and the dilation rate of 3 Conv layers of each branch are [16, 32, 64] and [(1,1), (2,1), (3,1)]. The dilation rate only changes along the time axis.

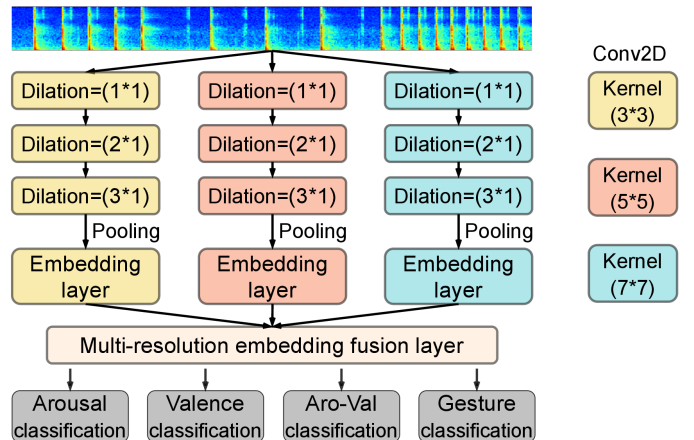


Fig. 4: The proposed lightweight multi-temporal resolution convolutional neural network (MTRCNN).

Taking the branch with the largest kernel (7, 7) as an example, according to the convolution RFS calculation formula,

$$R_i = (R_{i-1} - 1) \times stride + k \quad (1)$$

where  $R_i$  is the  $i$ -th Conv layer’s RFS relative to the input feature map,  $R_0 = 1$ , the Conv step size  $stride$  defaults to 1, and  $k$  is the Conv kernel size. If there is no pooling operation, according to Eq. (1), the 1st Conv layer’s RFS on the time axis is  $R_1 = 7$ . For dilated Conv, the formula for the RFS is

$$R_i = (R_{i-1} - 1) \times stride + k + (k - 1)(r - 1) \quad (2)$$

where  $r$  is the dilation rate. For the 2nd Conv layer with dilation rate (2, 1), on the time axis,  $R_2$  is 20. For the 3rd Conv layer,  $R_3$  is 38 on the time axis. That is, without pooling, MTRCNN requires that the length of the input be at least 38 frames. With a frame hop of 10ms, the corresponding length of the input clip is at least 0.38s. Identifying emotions or gestures within 0.38 seconds is challenging, even for humans. Moreover, if pooling is not used, the number of model parameters and the computation load will increase. After comprehensive trade-offs, we add pooling operations to these Conv layers, resulting in a minimum input audio length of 1.10s.

After the three Conv layers, the representations are fed into the following 64D embedding layer to learn embeddings with different resolutions. Then, the 3 branches' embeddings are concatenated and fed into the 192-dimensional ( $64 * 3$ ) multi-resolution embedding fusion layer to fuse information from different temporal resolutions. The fused embeddings are fed to the last classification layers. For the resulting MTRCNN model, the FLOPs is 0.708G, the number of parameters is 0.24M, the model size is 0.94MB, and it can process input audio clips of any length of at least 1.10s. These fit well with the design guidelines in Section (Sec.) III-A.

The tasks in subsequent experiments are all single-label multi-classification tasks, so the activation functions for the arousal (Aro), valence (Val), Aro-Val, and gesture classification layers are all Softmax, the loss function is cross entropy [32]. For more details, please see the *homepage* (<https://github.com/Yuanbo2020/MTRCNN>).

#### IV. EXPERIMENTS AND RESULTS

##### A. Dataset, experiments setup, and metrics

We conducted data collection experiments to record touch sounds for gesture and touch sounds for emotion [33]. As stated in Sec. II, participants first expressed gestures 10 emotions independently by touching Pepper's arm, and then expressed 6 touch gestures; the sounds generated by movements during touch are recorded as the dataset. Participants complete 3 rounds of interaction, each lasting 10s. Finally, there are 84 ( $28 * 3$ ) 10s audio clips for each gesture and emotion. For touch gesture classification, the number of samples in training, validation, and test sets is 366, 42, and 84, respectively. For emotion classification, the number of samples in training, validation, and test sets is 660, 80, and 100, respectively.

The Mel-filter with 64 banks is used as the acoustic feature, with a Hamming window of 32ms and an overlap of 10ms [34]. Dropout and normalization are used to prevent model overfitting [35]. A batch size of 32 and Adam optimizer [36] with a learning rate of 1e-3 are used to minimize loss. Models are trained on a GPU card Tesla T4 for 100 epochs, and 10 times without a fixed seed to obtain the mean result of 10 runs. Accuracy (Acc) is used to evaluate the classification results. Dataset, code, and models are available on the *homepage*.

##### B. Results and analysis

This part analyzes the performance of the proposed MTRCNN by the following research question (RQ).

**RQ1:** Is it feasible to recognise touch gestures and distinguish emotions' arousal and valence based on audio alone?

Table I shows the results of MTRCNN for classifying 6 gestures, as well as arousal and valence of emotions. Arousal is usually classified as low, neutral, and high. Valence is classified as negative, neutral, and positive. The AVDM of arousal-valence joint classification has four quadrants and an origin, so the arousal-valence [23] joint classification has five categories: Q1, Q2, Q3, Q4, and Q0, as shown in Fig. 3.

TABLE I: Test set classification results of the model with 10s audio clip input; models are repeated 10 times.

	Arousal (Aro)	Valence (Val)	Aro-Val	Gesture
Acc.	69.97 $\pm$ 4.40	62.90 $\pm$ 4.78	53.93 $\pm$ 3.41	82.14 $\pm$ 3.46

In Table I, MTRCNN performs better on touch gesture classification than emotion classification. This may be because gestures usually contain clear and consistent patterns, such as regular sounds when tapping and snapping sounds when patting, so the model can capture sounds produced by specific movements and rhythms to identify touch gestures effectively. However, due to the variety of ways in which different participants express emotions and different perceptions of the same type of emotions [37], e.g., a calm emotion expressed by some may appear sad to others. This makes Aro-Val-based emotion classification, especially relying on sounds produced when touching, challenging.

In Table I, MTRCNN has a higher classification accuracy on the arousal dimension than the valence dimension. Arousal denotes the intensity of emotion, which is usually conveyed via direct physical cues, e.g., pressure, frequency, and speed, which allows MTRCNN to grab these cues in sounds to efficiently identify the class of arousal. Valence reflects emotion's positive or negative nature, which is more subtle and context-dependent [38], making it challenging to distinguish it based solely on touch actions and the sounds caused by it.

**RQ2:** What is the shortest effective audio length required for touch-sound-based gesture and emotion recognition models?

Table I shows the performance of MTRCNN trained with full 10s audio clips. Here, we further explore the minimum audio length required for MTRCNN to effectively recognise gestures and emotion states. The input audio length can be regarded as a hyperparameter of the model. To avoid information leakage, Table II shows the results of this hyperparameter on the validation set. As mentioned in Sec. III-B, the proposed MTRCNN can handle audio clips with varying lengths with a minimum length of 1.10s, so the input audio length range in Table II is [1.10, 10].

Table II shows that the accuracy of touch gesture classification increases with the input audio length and peaks at 6s. The gesture classification results closest to the 6s result are at 5s and 7s, respectively, so we conduct statistical analysis on these similar results. The Shapiro-Wilk test [39] shows that the data follow a normal distribution. Then, the paired t-test [40] is used, and the statistics show that the gesture classification results based on 6s clips are significantly better than those of 5s

TABLE II: Accuracy of Aro-Val joint classification and gesture classification with varying input lengths on the validation set

Input length	1.10 s	2.00 s	3.00 s	4.00 s	5.00 s	6.00 s	7.00 s	8.00 s	9.00 s	10.00 s
Aro-Val	41.17±2.61	50.77±2.81	51.04±2.12	52.63±4.86	53.29±4.26	54.46±3.71	<b>58.54±3.74</b>	56.58±5.62	53.83±3.46	51.79±3.84
Gesture	68.81±3.06	84.99±2.76	85.48±2.37	85.71±3.17	88.47±2.46	<b>90.24±3.80</b>	86.19±3.33	83.81±3.69	83.57±6.39	79.31±6.55

( $t = 3.19, p < 0.05$ ) and  $7s$  ( $t = 3.22, p < 0.05$ ). In addition, for the Aro-Val joint classification of emotions, the results closest to  $7s$  are the results of  $6s$  and  $8s$ , respectively. Paired t-test for Aro-Val joint classification shows that the results based on the  $7s$  clip are significantly better than those of  $6s$  ( $t = 11.31, p < 0.05$ ) and  $8s$  ( $t = 2.58, p < 0.05$ ).

The above analyses show that MTRCNN effectively recognises different touch gestures within  $6s$  and decodes emotions within  $7s$ . Hence, input lengths of  $6s$  and  $7s$  will be used as default settings of MTRCNN for touch gesture classification and Aro-Val classification of emotions, respectively.

**RQ3:** What are the most challenging touch gestures and the emotions’ dimensions to distinguish based on sounds?

In Fig. 5 (a), MTRCNN performs better in identifying high arousal than low and neutral states. This implies that high arousal associated with touch sounds is easier to distinguish. Fig. 5 (c) implies that touch-based emotions in positive valence are more distinguishable than those in negative valence. In the Aro-Val space in Fig. 5 (b), MTRCNN can better distinguish emotions in Q2 (high arousal, negative valence) and Q4 (low arousal, positive valence) than those in Q1 and Q3. This is interesting because it suggests that these combinations of arousal and valence may be more consistently expressed by specific touch gestures conveyed by participants. The high and low arousal in Fig. 5 (a), as well as positive and negative valence in Fig. 5 (c), are rarely misclassified as neutral, implying that non-neutral emotions are less likely to be confused with neutral emotion due to their different tactile cues.

For gesture recognition, as shown in Fig. 5 (d), MTRCNN accurately identifies gestures such as *pat*, *tap*, and *hold*. These gestures usually have unique tactile profiles that are easy to identify. Like, *pat* may involve an easily recognizable repetitive rhythmic pattern [41], while *tap* and *hold* are simple, discrete movements with clear tactile features. In addition, *rub* is only misclassified as *hold*. This may be due to the subtle differences and overlapping tactile sensations between the two gestures, making their sounds similar, especially when the strength or speed of *rub* is not obvious.

**RQ4:** How does the proposed model perform compared to other typical sound-related models?

Table III compares models’ performance on the same server (CPU: AMD EPYC 7352, GPU: Tesla T4 16GB). CNN-Transformer consists of 3 convolutional layers with  $(3 \times 3)$  kernels, a Transformer encoder, and the final classification layers. YAMNet and MobileNetV2 [42] are classic CNN-based networks. PANNs [34] have shown excellent performance on AudioSet [43] and audio pattern-related tasks. Therefore, #4 and #5 explore the performance of PANNs with and without pretrained weights on large-scale AudioSet [43], respectively. The #5 with pretrained weights (PreW) significantly out-




Fig. 5: Normalized confusion matrix on the test set.

performs #4, indicating that the knowledge from the 5800-hour AudioSet effectively improves the model performance. However, #6, which has the smallest number of parameters and model size, achieves the best results in gesture classification, and is better than PANNs with PreW in #5. For Aro-Val classification, although #5, which uses PreW from AudioSet, is slightly better than the proposed #6 on the mean values, there is no statistically significant difference ( $t = -0.70, p > 0.05$ ).

TABLE III: Comparison of different models on the test set.

#	Model	Param. (M)	Size (MB)	FLOPs (G)	Inference time (s)	Accuracy	
						Aro-Val	Gesture
1	CNN-Trans.	1.58	6.02	0.266	0.006	28.87±4.47	60.60±6.51
2	YAMNet	3.21	12.30	0.728	0.008	29.03±2.77	61.19±5.78
3	MobileNetV2	2.23	8.74	0.351	0.007	45.63±6.22	71.90±5.56
4	PANNs	79.68	304.1	11.96	0.012	49.20±3.02	76.43±4.48
5	PANNs PreW	79.68	304.1	11.96	0.012	55.83±4.84	83.33±3.12
6	MTRCNN	<b>0.24</b>	<b>0.94</b>	0.708	0.007	54.73±3.29	<b>84.17±3.89</b>

## V. CONCLUSION

This paper explores the feasibility of identifying touch gestures and the emotions leading to them based on sounds produced by movements during touch, which fills the gap in HRI that lacks touch-related sounds to decode touch gestures and emotions. The proposed sound-based touch gesture and emotion recognition model can effectively recognize the arousal and valence states of different emotions, as well as various tactile gestures, when the input audio length varies from at least about 2 seconds to the optimal 6 to 7 seconds. Moreover, its lightweight, low-parameter, and low-latency processing characteristics make it ideal for real-time applications on robots such as Pepper. Future work will package the proposed model into an application and deploy it on Pepper.

## REFERENCES

- [1] H. Su, W. Qi, J. Chen, C. Yang, J. Sandoval, and M. A. Laribi, "Recent advancements in multimodal human-robot interaction," *Frontiers in Neurorobotics*, vol. 17, pp. 1084000, 2023.
- [2] M. Spezialetti, G. Placidi, and S. Rossi, "Emotion recognition for human-robot interaction: Recent advances and future perspectives," *Frontiers in Robotics and AI*, vol. 7, pp. 532279, 2020.
- [3] S. Marques-Villarroya, J. J. Gamboa-Montero, C. Jumela-Yedra, J. C. Castillo, and M. A. Salichs, "Affect display recognition through tactile and visual stimuli in a social robot," in *International Conference on Social Robotics*. Springer, 2022, pp. 130–140.
- [4] V. Villani, B. Capelli, C. Secchi, C. Fantuzzi, and L. Sabatini, "Humans interacting with multi-robot systems: a natural affect-based approach," *Autonomous Robots*, vol. 44, pp. 601–616, 2020.
- [5] L. Brown and A. M. Howard, "Gestural behavioral implementation on a humanoid robotic platform for effective social interaction," in *IEEE International Symposium on Robot and Human Interactive Communication*, 2014, pp. 471–476.
- [6] C. Basdogan, F. Giraud, V. Levesque, and S. Choi, "A review of surface haptics: Enabling tactile effects on touch surfaces," *IEEE Transactions on haptics*, vol. 13, no. 3, pp. 450–470, 2020.
- [7] A. K. Pandey, R. Gelin, and A. Robot, "Pepper: The first machine of its kind," *IEEE Robotics & Automation Magazine*, vol. 25, no. 3, pp. 40–48, 2018.
- [8] S. Shamsuddin, L. I. Ismail, H. Yussof, N. I. Zahari, S. Bahari, H. Hashim, and A. Jaffar, "Humanoid robot nao: Review of control and motion exploration," in *IEEE international conference on Control System, Computing and Engineering*. IEEE, 2011, pp. 511–516.
- [9] S. Al Moubayed, J. Beskow, G. Skantze, et al., "Furhat: a back-projected human-like robot head for multiparty human-machine interaction," in *Cognitive Behavioural Systems*. Springer, 2012, pp. 114–130.
- [10] R. S. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile sensing—from humans to humanoids," *IEEE Transactions on robotics*, vol. 26, no. 1, pp. 1–20, 2009.
- [11] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, pp. 2762, 2017.
- [12] M. Meribout, N. A. Takele, O. Derege, N. Rifiki, M. El K., V. Tiwari, and J. Zhong, "Tactile sensors: A review," *Measurement*, p. 115332, 2024.
- [13] Z. Ren, J. Meng, and J. Yuan, "Depth camera based hand gesture recognition and its applications in human-computer-interaction," in *International conference on information, communications & signal processing*. IEEE, 2011, pp. 1–5.
- [14] Z. Xia, Q. Lei, Y. Yang, H. Zhang, Y. He, W. Wang, and M. Huang, "Vision-based hand gesture recognition for human-robot collaboration: a survey," in *International Conference on Control, Automation and Robotics*. IEEE, 2019, pp. 198–205.
- [15] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake, "Real-time human pose recognition in parts from single depth images," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 2011, pp. 1297–1304.
- [16] F. Noroozi, C. A. Corneanu, D. Kamińska, T. Sapiński, S. Escalera, and G. Anbarjafari, "Survey on emotional body gesture recognition," *IEEE Transactions on affective computing*, vol. 12, no. 2, pp. 505–523, 2018.
- [17] K. Sun, T. Zhao, W. Wang, and L. Xie, "Vskin: Sensing touch gestures on surfaces of mobile devices using acoustic signals," in *Proceedings of Annual International Conference on Mobile Computing and Networking*, 2018, pp. 591–605.
- [18] A. Tajadura-Jiménez, N. Bianchi-Berthouze, E. Furfaro, and F. Bevilacqua, "Sonification of surface tapping changes behavior, surface perception, and emotion," *IEEE MultiMedia*, vol. 22, no. 1, pp. 48–57, 2015.
- [19] R. Cowie, E. Douglas-Cowie, N. Tsapatsoulis, G. Votsis, S. Kollias, et al., "Emotion recognition in human-computer interaction," *IEEE Signal processing magazine*, vol. 18, no. 1, pp. 32–80, 2001.
- [20] F. Eyben, M. Wöllmer, and B. Schuller, "Openear—introducing the munich open-source emotion and affect recognition toolkit," in *International conference on affective computing and intelligent interaction and workshops*. IEEE, 2009, pp. 1–6.
- [21] N. Gasteiger, J. Lim, M. Hellou, B. A. MacDonald, and H. S. Ahn, "A scoping review of the literature on prosodic elements related to emotional speech in human-robot interaction," *International Journal of Social Robotics*, vol. 16, no. 4, pp. 659–670, 2024.
- [22] J. A. Russell, M. Lewicka, and T. Niit, "A cross-cultural study of a circumplex model of affect," *Journal of personality and social psychology*, vol. 57, no. 5, pp. 848, 1989.
- [23] J. A. Russell, "A circumplex model of affect," *Journal of personality and social psychology*, vol. 39, no. 6, pp. 1161, 1980.
- [24] F. M. Citron, M. A. Gray, H. D. Critchley, B. S. Weekes, and E. C. Ferstl, "Emotional valence and arousal affect reading in an interactive way: neuroimaging evidence for an approach-withdrawal framework," *Neuropsychologia*, vol. 56, pp. 79–89, 2014.
- [25] T. Colibazzi, J. Posner, Z. Wang, D. Gorman, A. Gerber, S. Yu, et al., "Neural systems subserving valence and arousal during the experience of induced emotions," *Emotion*, vol. 10, no. 3, pp. 377, 2010.
- [26] D. Botteldooren, C. Lavandier, A. Preis, D. Dubois, I. Aspuru, et al., "Understanding urban and natural soundscapes," in *Forum Acusticum 2011*. European Acoustics Association, 2011, pp. 2047–2052.
- [27] D. Botteldooren, T. Andringa, I. Aspuru, L. Brown, D. Dubois, C. Guastavino, C. Lavandier, M. Nilsson, and A. Preis, "Soundscape for european cities and landscape: understanding and exchanging," in *COST TD0804 Final conference: Soundscape of European cities and landscapes*. Soundscape-COST, 2013, pp. 36–43.
- [28] Y. Hou, Q. Ren, H. Zhang, A. Mitchell, F. Aletta, J. Kang, and D. Botteldooren, "Ai-based soundscape analysis: Jointly identifying sound sources and predicting annoyance," *The Journal of the Acoustical Society of America (JASA)*, vol. 154, no. 5, pp. 3145–3157, 2023.
- [29] Y. Hou, Q. Ren, A. Mitchell, W. Wang, J. Kang, T. Belpaeme, and D. Botteldooren, "Soundscape captioning using sound affective quality network and large language model," *arXiv preprint arXiv:2406.05914*, 2024.
- [30] P. Wang, P. Chen, Y. Yuan, D. Liu, Z. Huang, X. Hou, and G. Cottrell, "Understanding convolution for semantic segmentation," in *IEEE winter conference on applications of computer vision*. IEEE, 2018, pp. 1451–1460.
- [31] H. Gao, Z. Chen, and C. Li, "Hierarchical shrinkage multiscale network for hyperspectral image classification with hierarchical feature fusion," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 5760–5772, 2021.
- [32] Y. Hou, B. Kang, A. Mitchell, W. Wang, J. Kang, and D. Botteldooren, "Cooperative scene-event modelling for acoustic scene classification," *IEEE/ACM Transactions on audio, speech, and language processing*, vol. 32, pp. 68–82, 2024.
- [33] Q. Ren, R. Proesmans, F. Bossuyt, J. Vanfleteren, F. Wyffels, and T. Belpaeme, "Conveying emotions to robots through touch and sound," *arXiv preprint arXiv:2412.03300*, 2024.
- [34] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, "PANNs: Large-scale pretrained audio neural networks for audio pattern recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [35] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, et al., "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. of ICLR*, 2015.
- [37] J. R. Zadra and G. L. Clore, "Emotion and perception: The role of affective information," *Wiley interdisciplinary reviews: cognitive science*, vol. 2, no. 6, pp. 676–685, 2011.
- [38] B. De Martino, "The effect of context on choice and value," in *Neuroscience of Preference and Choice*, pp. 93–119. Elsevier, 2012.
- [39] B. W. Yap and C. H. Sim, "Comparisons of various types of normality tests," *Journal of Statistical Computation and Simulation*, vol. 81, no. 12, pp. 2141–2155, 2011.
- [40] M. D. Riina, C. Stambaugh, N. Stambaugh, and K. E. Huber, "Continuous variable analyses: t-test, mann-whitney, wilcoxon rank," in *Translational Radiation Oncology*, pp. 153–163. Elsevier, 2023.
- [41] N. Hogan and D. Sternad, "On rhythmic and discrete movements: reflections, definitions and implications for motor control," *Experimental brain research*, vol. 181, pp. 13–30, 2007.
- [42] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4510–4520.
- [43] J. Gemmeke, D. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, et al., "AudioSet: An ontology and human-labeled dataset for audio events," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.