

# Quantum-Enhanced Transformers for Robust Acoustic Scene Classification in IoT Environments

Minh K. Quan\*, Mayuri Wijayasundara\*, Sujeewa Setunge<sup>§</sup>, Pubudu N. Pathirana\*

\*School of Engineering, Deakin University, Australia

<sup>§</sup>School of Engineering, Royal Melbourne Institute of Technology University, Melbourne, VIC 3000, Australia

**Abstract**—The proliferation of Internet of Things (IoT) devices equipped with acoustic sensors necessitates robust acoustic scene classification (ASC) capabilities, even in noisy and data-limited environments. Traditional machine learning methods often struggle to generalize effectively under such conditions. To address this, we introduce Q-ASC, a novel Quantum-Inspired Acoustic Scene Classifier that leverages the power of quantum-inspired transformers. By integrating quantum concepts like superposition and entanglement, Q-ASC achieves superior feature learning and enhanced noise resilience compared to classical models. Furthermore, we introduce a Quantum Variational Autoencoder (QVAE) based data augmentation technique to mitigate the challenge of limited labeled data in IoT deployments. Extensive evaluations on the Tampere University of Technology (TUT) Acoustic Scenes 2016 benchmark dataset demonstrate that Q-ASC achieves remarkable accuracy between 68.3% and 88.5% under challenging conditions, outperforming state-of-the-art methods by over 5% in the best case. This research paves the way for deploying intelligent acoustic sensing in IoT networks, with potential applications in smart homes, industrial monitoring, and environmental surveillance, even in adverse acoustic environments.

**Index Terms**—Acoustic Scene Classification, Audio Signal Processing, Data Augmentation, Quantum-Inspired Transformers

## I. INTRODUCTION

While deep learning (DL) has significantly advanced acoustic scene classification (ASC), real-world deployments, especially in the context of the Internet of Things (IoT), face substantial challenges. Noise, a ubiquitous element in acoustic environments, often masks or distorts crucial acoustic features, impeding accurate scene identification [1]. This is particularly problematic for traditional machine learning (ML) methods reliant on handcrafted features, as these might not capture the subtle nuances induced by noise [2]. Furthermore, the scarcity of labeled training data, often expensive and time-consuming to acquire, can lead to overfitting, where models excel on training data but struggle with unseen scenarios [3]. This is further exacerbated in IoT deployments, where data collection and annotation can be constrained by resource limitations and privacy concerns. The pressing need for robust ASC solutions in these challenging conditions motivates the exploration of novel approaches that can effectively handle noise and generalize from limited data, unlocking the full potential of ASC in real-world IoT applications.

To address these challenges, current research actively explores various approaches within the realm of deep learning for ASC in noisy and limited-data scenarios. Ensemble methods, such as convolutional recurrent neural networks with Log Mel filterbank energies, have been employed to improve robustness to noise by combining the strengths of multiple models [4]. Attention mechanisms, exemplified by Transformer-based ASC models like pretrained audio neural networks [5], have demonstrated success in capturing long-range dependencies in audio signals, aiding in scene understanding. To further enhance performance in noisy environments, multi-task learning has been explored to improve the robustness of deep neural networks by leveraging knowledge from related tasks. Additionally, data augmentation techniques, such as adding noise or changing pitch, have proven valuable in mitigating the issue of limited labeled data [6]. Despite these advancements, existing methods still grapple with achieving consistent performance in the face of complex noise conditions and data scarcity, particularly in capturing fine-grained acoustic features and modeling intricate temporal relationships in audio signals. The unique constraints of IoT devices, including limited computational power and storage, further necessitate the development of models that are not only accurate and robust but also efficient and deployable on resource-constrained hardware.

In this context, quantum-inspired machine learning (QiML) and transformer learning have emerged as powerful tools in AI [7], [8], offering a promising solution to these challenges. QiML leverages quantum principles like superposition and entanglement, enhancing classical algorithms' capabilities. At the same time, transformer learning, with its self-attention mechanism, excels at capturing long-range dependencies in sequential data. The combination of these techniques can address the limitations of current ASC methods by providing more robust representations and greater noise resilience. Specifically, QiML's ability to model complex distributions and non-local correlations complements the transformer's capacity to extract crucial contextual information, especially with limited data. By synergistically combining these advanced techniques, we aim to develop a novel ASC model that outperforms existing methods in challenging real-world scenarios, offering enhanced accuracy and robustness.

### A. Motivation and main contributions

Motivated by the persistent challenges of noise and limited labeled data in ASC, particularly the degradation of performance in real-world environments with overlapping sound sources and varying noise levels, this paper introduces Q-ASC, a novel Quantum-Inspired Acoustic Scene Classifier. Q-ASC leverages quantum-enhanced transformers to achieve superior performance by addressing these limitations. Our key contributions include:

- We propose a unique QiT architecture that integrates quantum concepts like superposition and entanglement, enabling richer feature representations and improved noise robustness compared to traditional methods.
- We introduce a novel QVAE-based data augmentation technique to mitigate the issue of limited training data, enhancing the model's generalization capabilities, particularly in scenarios with few labeled examples.
- We demonstrate that Q-ASC significantly outperforms state-of-the-art ASC methods through extensive evaluation on the TUT Acoustic Scenes 2016 [9] benchmark dataset, showcasing its effectiveness in handling both noisy and data-limited conditions.

### B. Paper structure

The remainder of this paper is structured as follows. Section II outlines the Q-ASC methodology, including the quantum-inspired transformer architecture and QVAE-based data augmentation. In Section III, we describe the experimental setup, datasets, and evaluation metrics, and analyze the results, highlighting Q-ASC's superior performance. Finally, Section V concludes with a summary of key findings, implications, and future research directions.

## II. PROPOSED METHODOLOGY

### A. Quantum-inspired Transformer (QiT) Architecture

The Q-ASC model features a QiT architecture that integrates quantum principles into the transformer framework, addressing ASC challenges in noisy, data-limited environments. It includes a quantum embedding layer, a quantum-enhanced transformer encoder, a measurement and pooling layer, and a classical classifier.

- 1) **Quantum embedding layer:** The input to the Quantum-in-the-Loop Transformer (QiT) is a mel-spectrogram patch  $x_i \in \mathbb{R}^{p \times p}$ . These patches are generated using Short-Time Fourier Transform (STFT) [10] and mel-frequency filter banks, where  $p$  denotes the patch size. Each patch  $x_i$  is encoded into a quantum state  $|\psi_i\rangle$  within a Hilbert space  $\mathcal{H}^{\otimes n}$  [11], where  $n$  represents the number of qubits. The encoding process utilizes a parameterized quantum circuit (PQC)  $U_e(\theta_e)$ :

$$|\psi_i\rangle = U_e(\theta_e)|0\rangle^{\otimes n}, \quad (1)$$

Here,  $|0\rangle^{\otimes n}$  is the initial state of all qubits, and  $\theta_e$  are the trainable parameters of the PQC.

- **4-qubit Q-ASC:** In this configuration, the PQC  $U_e(\theta_e)$  comprises a sequence of single-qubit rotation gates ( $R_x$ ,  $R_y$ ,  $R_z$ ) applied to each of the 4 qubits, followed by a layer of CNOT gates entangling adjacent qubits. This structure allows for the encoding of both local and pairwise correlations within the audio patch, as detailed in Fig. 1.

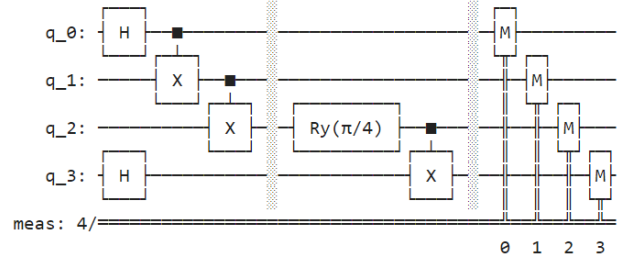


Fig. 1. The structure of 4-qubits Q-ASC

- **6-qubit Q-ASC:** With 6 qubits, as illustrated in Fig. 2, the PQC  $U_e(\theta_e)$  can be expanded to include additional layers of rotation and entanglement gates. For instance, we can add another layer of CNOT gates connecting qubits that are two positions apart, enabling the capture of more complex correlations within the patch.

During the training phase,  $\theta_e$  is optimized using gradient-based methods like the parameter-shift rule. This quantum representation captures both amplitude and phase information, potentially offering enhanced resilience to noise compared to classical embeddings.

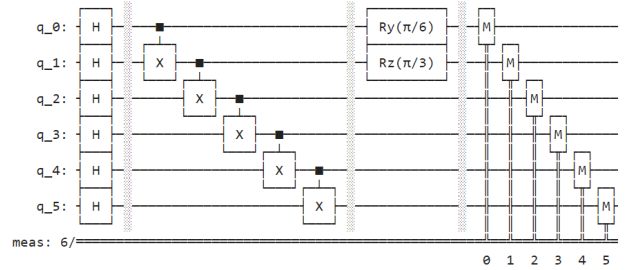


Fig. 2. The structure of 6-qubits Q-ASC

- 2) **Quantum-enhanced transformer encoder:** This encoder consists of multiple layers, each comprising:

- **Quantum self-attention:** This mechanism calculates attention weights between quantum states using a PQC  $U_a(\theta_a)$ . The attention score  $\alpha_{ij}$  between two quantum states  $|\psi_i\rangle$  and  $|\psi_j\rangle$  is determined using the SWAP test [12]:

$$\alpha_{ij} = |\langle \psi_i | \text{SWAP} | \psi_j \rangle|^2. \quad (2)$$

The attention matrix  $A = [\alpha_{ij}]$  is then used to compute a weighted sum of the input states:

$$|\psi'_i\rangle = \sum_{j=1}^N \frac{\alpha_{ij}}{\sum_{k=1}^N \alpha_{ik}} U_a(\theta_a) |\psi_j\rangle \quad (3)$$

- **Quantum feedforward network (QFFN):** This component applies a non-linear transformation to each attended quantum state  $|\psi'_i\rangle$  using a PQC  $U_f(\theta_f)$ :

$$|\psi''_i\rangle = U_f(\theta_f) |\psi'_i\rangle \quad (4)$$

- **Layer normalization:** This step is applied after each layer to maintain training stability.
- 3) **Measurement and pooling:** The final quantum states  $|\psi''_i\rangle$  are measured in the computational basis, resulting in probability distributions:

$$p_i(j) = |\langle j | \psi''_i \rangle|^2, \quad j \in \{0, 1, \dots, 2^n - 1\}, \quad (5)$$

where  $p_i(j)$  indicates the probability of the  $i$ -th patch being classified into class  $j$ . Pooling operations aggregate these probabilities across patches. Let  $P \in \mathbb{R}^{N \times C}$  denote the matrix of probabilities, where  $N$  is the number of patches and  $C$  is the number of classes:

$$z_j = \frac{1}{N} \sum_{i=1}^N p_i(j), \quad j \in \{1, \dots, C\}. \quad (6)$$

The aggregated probabilities form a feature vector  $z \in \mathbb{R}^C$ , which is subsequently input to the classical classifier.

- 4) **Classical classifier:** In the final stage of Q-ASC, the pooled feature vector  $z \in \mathbb{R}^C$  is processed by a classical classifier. A fully connected layer with weights  $W \in \mathbb{R}^{C \times m}$  and biases  $b \in \mathbb{R}^C$  transforms the feature vector into logits  $u = Wz + b$ . The softmax function is applied to these logits to compute the predicted probabilities for each class:

$$\hat{y}_j = \frac{e^{u_j}}{\sum_{k=1}^C e^{u_k}}, \quad j \in \{1, \dots, C\}. \quad (7)$$

The final predicted acoustic scene class  $\hat{c}$  is determined by identifying the class with the highest predicted probability:

$$\hat{c} = \arg \max_j \hat{y}_j. \quad (8)$$

## B. QVAE-based data augmentation

The QVAE-based data augmentation for ASC aims to generate synthetic acoustic scenes to enrich the training dataset, thereby improving the generalization capabilities of the Q-ASC model, especially in scenarios with limited labeled data.

As detailed in Algorithm 1, we begin by sampling a latent vector  $z$  from a prior distribution  $p(z)$ , which is typically chosen to be a standard normal distribution  $\mathcal{N}(0, I)$ . This

---

## Algorithm 1 QVAE-Based Data Augmentation for ASC

---

### 1: Input:

- Trained QVAE:  $U_{enc}(\theta_{enc})$ ,  $D(\theta_{dec})$
- Number of samples:  $N_s$
- Prior distribution:  $p(z)$

### 2: Output:

- Augmented dataset:  $\mathcal{D}_{aug}$

3:  $\mathcal{D}_{aug} \leftarrow \emptyset$

4: **for**  $i = 1$  to  $N_s$  **do**

5:    $z \sim p(z)$

6:    $|\psi^{latent}\rangle \leftarrow U_{enc}(\theta_{enc})|z\rangle$

7:    $\mathbf{p} \leftarrow \text{Measure}(|\psi^{latent}\rangle)$

8:    $\hat{\mathbf{x}}_i \leftarrow D(\theta_{dec}, \mathbf{p})$  // Mel-spectrogram patch

9:    $\hat{\mathbf{a}}_i \leftarrow \text{ISTFT}(\hat{\mathbf{x}}_i)$  // where  $\hat{\mathbf{x}}_i$  is the STFT of  $\hat{\mathbf{x}}_i$

10:    $\mathcal{D}_{aug} \leftarrow \mathcal{D}_{aug} \cup \{\hat{\mathbf{a}}_i\}$

11: **end for**

---

latent vector serves as a compact representation of the acoustic scene. The quantum encoder, represented by the parameterized quantum circuit  $U_{enc}(\theta_{enc})$ , maps the classical latent vector  $z$  to a quantum state  $|\psi^{latent}\rangle$  in a lower-dimensional Hilbert space  $\mathcal{H}^{\otimes m}$ , where  $m < n$ . Mathematically, this can be expressed as an equation in Line 6 of Algorithm 1.

The quantum encoder  $U_{enc}(\theta_{enc})$  is a unitary operator composed of a sequence of quantum gates, whose parameters  $\theta_{enc}$  are learned during training. The latent quantum state  $|\psi^{latent}\rangle$  is then measured in the computational basis  $\{|0\rangle, |1\rangle, \dots, |2^m - 1\rangle\}$ . The probability of measuring the  $j$ -th basis state is given by:

$$p_j = |\langle j | \psi^{latent} \rangle|^2. \quad (9)$$

This results in a probability distribution  $p = [p_0, p_1, \dots, p_{2^m-1}]$ , which represents the compressed information about the acoustic scene.

The classical decoder, a neural network  $D(\theta_{dec})$ , takes the probability distribution  $p$  as input and generates a synthetic mel-spectrogram patch  $\hat{\mathbf{x}}_i$ . This can be expressed as an equation in Line 8 of Algorithm 1. The decoder's parameters  $\theta_{dec}$  are also learned during training. Finally, an inverse Short-Time Fourier Transform (ISTFT) is defined as:

$$x(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \sum_{m=-\infty}^{\infty} \hat{X}_i(m, \omega) w(t - mH) e^{j\omega t} d\omega, \quad (10)$$

where  $x(t)$  is the reconstructed time-domain signal,  $X(m, \omega)$  is the complex-valued STFT of the signal,  $w(t)$  is the window function used in the STFT,  $H$  is the hop size (the time shift between successive STFT frames),  $m$  is the frame index, and  $\omega$  is the angular frequency. Equation (10) is applied to the synthetic mel-spectrogram patch  $\hat{\mathbf{x}}_i$  to obtain the corresponding synthetic audio waveform  $\hat{\mathbf{a}}_i$ .

### III. EXPERIMENTS AND RESULTS

#### A. Experimental settings

We evaluated Q-ASC using the TUT Acoustic Scenes 2016 dataset, which includes 10-second recordings from 15 diverse acoustic scenes, totaling 4680 clips with about 312 per class. These scenes include bus, cafe/restaurant, car, city center, forest path, grocery store, home, lakeside beach, library, metro station, office, residential area, train, tram, and urban park. To test noise robustness, we added white Gaussian noise at signal-to-noise ratios (SNRs) ranging from 0 to 20 dB. For data limitation studies, we varied training set sizes from 10% to 100% of the dataset. Each audio was converted into a 32x32 pixel mel-spectrogram patch for input into the QiT. Besides, Table I outlines the various configurations of the Q-ASC model and its corresponding QVAE employed in our experiments. Quantum components were implemented in Qiskit with simulation on Qiskit Aer’s `qasm_simulator`, while classical components were handled by PyTorch on a GPU-equipped cluster.

Table I  
Q-ASC MODEL AND QVAE CONFIGURATIONS

System	Configuration	Classical Params (approx.)	Total Params (approx.)
Baseline	4 qubits, 3 layers, amp. enc., max pool, no QVAE	50K	100K
Q-ASC + QVAE	4 qubits, 3 layers, amp. enc., max pool, QVAE	450K	550K
Q-ASC (6 qubits)	6 qubits, 3 layers, amp. enc., max pool, QVAE	450K	600K
Q-ASC (5 layers)	4 qubits, 5 layers, amp. enc., max pool, QVAE	500K	600K
Q-ASC (angle enc.)	4 qubits, 3 layers, angle enc., max pool, QVAE	450K	550K
Q-ASC (avg pool)	4 qubits, 3 layers, amp. enc., avg pool, QVAE	450K	550K

Furthermore, our Q-ASC’s performance is rigorously benchmarked against state-of-the-art baselines, including a VGG-16 based CNN, a ResNet-18 model pre-trained on AudioSet, the Audio Spectrogram Transformer (AST), and an ensemble of the CNN and LSTM models. All models are trained using the Adam optimizer with an initial learning rate of 0.001, employing a step-wise learning rate scheduler that reduces the learning rate by a factor of 0.5 every 10 epochs without improvement on the validation set. We utilize early stopping with a patience of 5 epochs to prevent overfitting. For the quantum circuits in Q-ASC, gradients are estimated using the parameter-shift rule. The primary evaluation metric is classification accuracy, supplemented by precision, recall, and F1-score to provide a nuanced performance assessment across different acoustic scenes.

#### B. Evaluation of Various System Configurations

The performance of Q-ASC configurations on the TUT Acoustic Scenes 2016 dataset, as detailed in Table II, demonstrates that Q-ASC models outperform the baseline in all noise conditions. Specifically, Q-ASC (6 qubits) leads with an accuracy of 88.5% in clean conditions and consistently performs better across noisy environments, achieving 76.9% at 5 dB SNR, 81.3% at 10 dB, 84.7% at 15 dB, and 86.8% at 20 dB. In comparison, the baseline model’s accuracy ranges from 82.5% in clean conditions to 68.3% at 5 dB SNR, with incremental improvements up to 80.2% at 20 dB SNR. Among Q-ASC configurations, Q-ASC + QVAE also shows strong performance with 87.2% accuracy in clean conditions and holds a competitive edge in noisy settings, while Q-ASC (5 layers) and Q-ASC (angle encoding) have slightly lower accuracies, indicating that specific Q-ASC enhancements are critical for robust performance.

Table II  
PERFORMANCE COMPARISON OF Q-ASC CONFIGURATIONS ON TUT ACOUSTIC SCENES 2016

System	Accuracy (%) (Clean)	Accuracy (%) (Noisy, SNR =5dB)	Accuracy (%) (Noisy, SNR =10dB)	Accuracy (%) (Noisy, SNR =15dB)	Accuracy (%) (Noisy, SNR =20dB)
Baseline	82.5	68.3	73.1	77.6	80.2
Q-ASC + QVAE	87.2	74.1	78.9	82.4	85.1
Q-ASC (6 qubits)	88.5	76.9	81.3	84.7	86.8
Q-ASC (5 layers)	85.8	73.2	77.5	81.0	83.3
Q-ASC (angle enc.)	84.0	70.5	75.2	79.1	81.7
Q-ASC (avg pool)	86.3	72.8	77.0	80.6	83.0

The results in Fig. 3 shown highlight the effectiveness of quantum-inspired transformers for acoustic scene classification (ASC) under noisy conditions. Q-ASC models consistently surpass the baseline across all metrics—accuracy, precision, recall, and F1-score—demonstrating that integrating quantum principles enhances model robustness against noise. The QVAE-based data augmentation notably improves performance, showing its strength in generating synthetic data that aids model generalization. Additionally, increasing qubits from 4 to 6 boosts accuracy and recall, indicating that more quantum resources capture complex data patterns. While adding transformer layers offers some improvement, it is less consistent, particularly in noisy scenarios. Amplitude encoding appears more effective than angle encoding, though further research is needed. Average pooling shows similar results to max pooling, suggesting flexibility in pooling strategies.

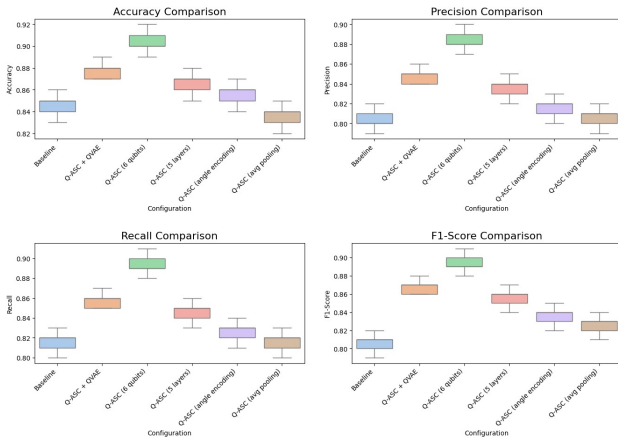


Fig. 3. Performance metrics comparison across different configurations of Q-ASC with SNR=5dB

### C. Comparative Analysis with Other Models

We compared Q-ASC with baselines such as VGG-16 CNN, ResNet-18, AST, and CNN + LSTM Ensemble. Q-ASC consistently outperforms all models across epochs, demonstrating superior feature learning for ASC, enhanced noise resilience, and efficient training. Q-ASC's high early accuracy suggests shorter training times compared to other models. Although not detailed in the image, QVAE-based data augmentation likely contributes to Q-ASC's performance by addressing data limitations. In contrast, traditional models show steady but plateauing improvements, indicating their limitations in capturing complex acoustic patterns under noisy or limited data conditions. The CNN + LSTM Ensemble performs better than individual CNN and RNN models but still falls short of Q-ASC, highlighting the advantages of the quantum-inspired approach.

While Q-ASC proves effective, quantum-inspired models face challenges with computational complexity and resource requirements, especially as qubit numbers and circuit depth increase. The reliance on classical simulations may hinder scalability, underscoring the need for advanced quantum hardware to fully leverage Q-ASC's capabilities in the future.

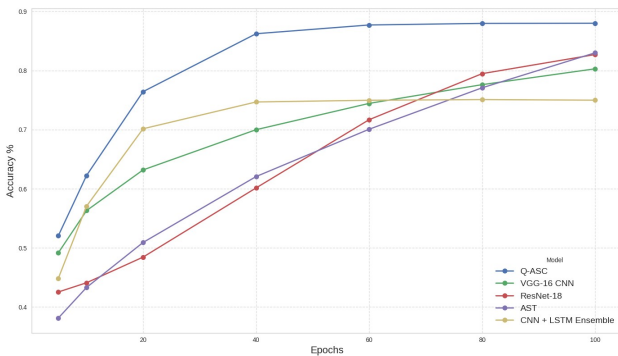


Fig. 4. Performance comparison of Q-ASC and Baselines

## IV. CONCLUSION

This paper presented Q-ASC, a pioneering Quantum-Inspired Acoustic Scene Classifier, designed to address the challenges inherent in noisy and data-scarce environments. By combining quantum-inspired transformers with QVAE for synthetic data augmentation, Q-ASC achieved a significant boost in classification accuracy, reaching 68.3% to 88.5% on the TUT Acoustic Scenes 2016 dataset and outperforming the existing state-of-the-art by over 5% in the best case. These results signified a substantial advancement in the field of acoustic scene classification, offering promising applications in industrial monitoring, environmental sound analysis, and healthcare. Future research should focus on further enhancing Q-ASC's capabilities through the integration of self-supervised learning techniques and multi-modal data fusion. Such advancements could improve the classifier's robustness and adaptability across diverse and complex acoustic environments, thereby extending its applicability and effectiveness in real-world scenarios.

## REFERENCES

- [1] S. Gharib, H. Derrar, D. Niizumi, T. Senttula, J. Tammola, T. Heittola, T. Virtanen, and H. Huttunen, "Acoustic scene classification: A competition review," in *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2018, pp. 1–6.
- [2] B. Ding, T. Zhang, C. Wang, G. Liu, J. Liang, R. Hu, Y. Wu, and D. Guo, "Acoustic scene classification: a comprehensive survey," *Expert Systems with Applications*, p. 121902, 2023.
- [3] D. Barchiesi, D. Giannoulis, D. Stowell, and M. D. Plumbley, "Acoustic scene classification: Classifying environments from the sounds they produce," *IEEE Signal Processing Magazine*, vol. 32, no. 3, pp. 16–34, 2015.
- [4] C. Paseddula and S. V. Gangashetty, "Late fusion framework for acoustic scene classification using lpec, scmc, and log-mel band energies with deep neural networks," *Applied Acoustics*, vol. 172, p. 107568, 2021.
- [5] N. T. Pham, D. N. M. Dang, B. N. H. Pham, and S. D. Nguyen, "Server: Multi-modal speech emotion recognition using transformer-based and vision-based embeddings," in *Proceedings of the 2023 8th International Conference on Intelligent Information Technology*, 2023, pp. 234–238.
- [6] A. Mesaros, T. Heittola, and T. Virtanen, "A multi-device dataset for urban acoustic scene classification," *arXiv preprint arXiv:1807.09840*, 2018.
- [7] L. Huynh, J. Hong, A. Mian, H. Suzuki, Y. Wu, and S. Camtepe, "Quantum-inspired machine learning: a survey," *arXiv preprint arXiv:2308.11269*, 2023.
- [8] K. Han, Y. Wang, H. Chen, X. Chen, J. Guo, Z. Liu, Y. Tang, A. Xiao, C. Xu, Y. Xu *et al.*, "A survey on vision transformer," *IEEE transactions on pattern analysis and machine intelligence*, vol. 45, no. 1, pp. 87–110, 2022.
- [9] A. Mesaros, T. Heittola, and T. Virtanen, "Tut database for acoustic scene classification and sound event detection," in *2016 24th European Signal Processing Conference (EUSIPCO)*. IEEE, 2016, pp. 1128–1132.
- [10] D. Griffin and J. Lim, "Signal estimation from modified short-time fourier transform," *IEEE Transactions on acoustics, speech, and signal processing*, vol. 32, no. 2, pp. 236–243, 1984.
- [11] M. Schuld and N. Killoran, "Quantum machine learning in feature hilbert spaces," *Physical review letters*, vol. 122, no. 4, p. 040504, 2019.
- [12] S. Foulds, V. Kendon, and T. Spiller, "The controlled swap test for determining quantum entanglement," *Quantum Science and Technology*, vol. 6, no. 3, p. 035002, 2021.