Data Science and Engineering (DSE) Record, Volume 6, issue 1.

REFRIGERANT LEAK DETECTION BY MACHINE LEARNING

Poompatai Muennamnor[1] and Pruet Boonma[2]

1 Master's Degree Program in Data Science, Chiang Mai University, Chiang Mai, Thailand
2 Department of Computer Engineering, Faculty of Engineering, Chiang Mai University, Chiang Mai, Thailand
poompatai_m@cmu.ac.th

Abstract. Refrigerant leaks from cooling systems can harm the environment and cost businesses money. Current ways to find leaks can be slow, expensive, and not always accurate. This project uses machine learning to create a better way to detect refrigerant leaks by listening to the sounds they make. The goal is to develop a system that can automatically and cheaply detect leaks early on, reducing environmental damage and saving businesses money. The system uses a microphone to record sounds, then a computer program analyzes the sounds to identify leaks. By using sound analysis, the system can tell the difference between normal sounds and the sounds of a refrigerant leak. This helps catch leaks early, lowers maintenance costs, and reduces greenhouse gas emissions.

Keywords: Refrigerant leaks, Cooling systems, Leak detection, Machine learning.

1 Introduction

1.1 Background and Problem Statement

Today, there are several methods available for detecting refrigerant leaks, including sensory detection techniques such as sight, smell, and hearing, as well as devices like thermal cameras and leak detectors. While these methods can be useful, they still have significant limitations. These approaches can be complex, slow in terms of operational efficiency, and costly, especially when applied on a larger scale. As a result, businesses and industries need a more efficient, cost-effective solution for detecting refrigerant leaks. Refrigerants, the substances used in cooling systems, are known to have a major impact on global warming. Many refrigerants contain chemicals that contribute to the depletion of the Earth's ozone layer, which in turn accelerates climate change. Therefore, it is crucial to detect and manage refrigerant leaks promptly to reduce environmental damage. In addition to the environmental consequences, refrigerant leaks can lead to significant financial losses. Leaking refrigerants not only affect the efficiency of cooling systems, but they can also increase energy consumption, further driving up costs. For industries such as hotels, food storage, and large-scale industrial operations, the cost of refrigerant leaks can be especially harmful to operations and profitability. Efficient detection and management of refrigerant leaks are essential to both protecting the environment and reducing financial losses for businesses, especially in sectors that rely heavily on refrigeration. The development of better detection methods is key to addressing these challenges.

1.2 Objectives

This project aims to develop a machine learning model for refrigerant leak detection to improve accuracy, reduce costs, and enhance efficiency. Traditional methods, such as manual inspection, leak detectors, and thermal cameras, can be expensive, time-consuming, and ineffective for continuous monitoring. Refrigerant leaks are a serious problem because they contribute to global warming, increase electricity costs, and cause equipment damage in industries like factories, hotels, cold storage, and supermarkets. This project uses machine learning and sound analysis to detect leaks by identifying patterns in ultrasonic sound waves, helping to catch leaks early and prevent financial and environmental damage. The goal is to create an automatic, affordable, and accurate system that improves how refrigerant leaks are detected in real-world conditions.

1.3 Outcomes

The outcomes of this project include the development of an accurate and efficient machine learning model for refrigerant leak detection using sound wave analysis. The model successfully differentiates between normal operational sounds and leak sounds, improving early detection and real-time monitoring. This approach reduces manual inspection efforts, lowers maintenance costs, and enhances energy efficiency for industries like factories, hotels, cold storage, and supermarkets. Additionally, the system helps reduce greenhouse gas emissions, contributing to environmental sustainability. By providing a cost-effective and automated solution, this work ensures better leak prevention and management, minimizing financial losses and operational risks for businesses.

2. Literature Review

2.1 Machine Learning for Leak Detection

The application of machine learning (ML) for fault diagnosis in industrial systems has gained significant attention in recent years, particularly in the context of Industry 4.0 (Dalzochio et al., 2020). This approach leverages data-driven models to detect anomalies and predict failures, offering a promising alternative to traditional, often costly, methods. In the specific domain of refrigerant leak detection, ML algorithms have shown potential for early and accurate identification of leaks, contributing to improved system efficiency, cost savings, and environmental protection (Mtibaa et al., 2024). In the context of fault detection, explainable ML models have been proposed to address the need for interpretable diagnostics. For example, Kim or author demonstrated the potential of convolutional neural networks (CNNs) in detecting faults in industrial motion systems, emphasizing the importance of model transparency. Similar approaches can be adapted for refrigerant leak diagnosis to build trust and ensure practical deployment.

2.2 Common Models and Feature Engineering Methods

In machine learning (ML) for refrigerant leak detection, the choice of model and the quality of features play a critical role in system performance. Various ML algorithms have been explored, including Artificial Neural Networks (ANNs), Support Vector Machines (SVMs), Bayesian Networks, and decision tree-based methods like Random Forests and Extremely Randomized Trees (Lei et al., 2022). In particular, tree-based models are often favored for their interpretability and strong performance with structured sensor data. Ensemble learning, which combines multiple models to improve prediction accuracy, has also shown strong results in fault detection tasks (Mtibaa et al., 2024).

Explainable models, such as Convolutional Neural Networks (CNNs), are gaining attention due to their ability to provide visual and transparent insights into how decisions are made. This is particularly important in critical systems where understanding the model's reasoning is just as important as the output itself (Lei et al., 2022).

Equally important is feature engineering—the process of converting raw sensor signals into meaningful inputs for ML models. Selecting the right features, such as pressure levels, temperature differences, and acoustic signal patterns, helps the model distinguish between normal and leak conditions (Kim et al., 2021). Effective feature selection ensures that the model remains accurate and efficient across different conditions. Studies have also shown that domain-specific features, when carefully designed, can significantly improve the robustness and reliability of leak detection systems (Mtibaa et al., 2024).

In this project, both feature engineering and model selection were tailored to focus on acoustic data, using frequency-based features extracted through FFT, STFT, and Wavelet transforms. A Decision Tree model was chosen for its interpretability and ease of use, making it suitable for early-stage testing and validation.

2.3 Limitations of Existing Approaches

While machine learning (ML) offers promising solutions for refrigerant leak detection, several limitations still exist in real-world applications. One major challenge is dealing with noisy environments, where background sounds can interfere with accurate leak detection. Incomplete or low-quality data from sensors can also affect model performance, especially in industrial settings where sensor reliability varies (Mtibaa et al., 2024). In addition, current models often struggle to distinguish between different types of system faults, such as leaks versus normal vibrations, making false positives or missed detections possible (Kim et al., 2021).

Another limitation is the difficulty in adapting ML models to different system configurations or operating conditions. A model trained on one machine may not perform well on another due to differences in system design or usage patterns. This limits the generalizability of many current approaches.

To overcome these issues, future work should focus on creating more robust and flexible models that can handle various data conditions and adapt to different environments. Researchers also suggest integrating ML with other technologies such as acoustic sensors and ontologies to enhance fault detection capabilities (Dalzochio et al., 2020). Additionally, the use of explainable AI (XAI) can improve user trust by making model decisions more transparent and understandable, especially in high-stakes applications like HVAC or industrial cooling systems.

3. Methodology

3.1 System Design Overview

From the flowchart, when a refrigerant leak occurs or there is suspicion of a leak, the first step is Data Mining or Data Collection. This is done using a microphone to record sound from the environment. The recorded audio is then analyzed and processed to prepare the data for the next stage. The prepared data is fed into a Machine Learning Model, where the model processes the sound features and determines whether the recorded sound indicates a leak or not. The final output of the system is a classification result: "Leak" or "Not Leak", which can be used for further decision-making and preventive measures.

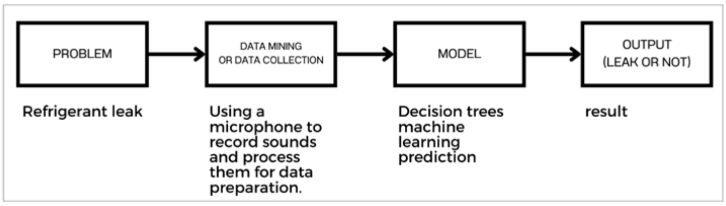

Figure 3.1 : Preliminary Model Workflow Diagram

3.2 Data Collection and Preparation

This project utilizes two types of audio data to train and evaluate the machine learning model for refrigerant leak detection: Generated Data and Real-World Data. Both types contribute to improving the model’s ability to recognize refrigerant leak sounds in a variety of conditions.

Generated data was created by first calculating expected leak frequencies based on known system parameters, such as pipe dimensions, flow characteristics, and refrigerant properties. Synthetic sounds were then generated at these target frequencies. To better mimic real-world conditions, these leak-like sounds were mixed with background environmental noises, including sounds from factory machinery, air compressors, and other industrial sources. This helps the model learn to distinguish true leak signals even in noisy environments.

In parallel, real-world data was collected through controlled laboratory experiments. A water cooler system using R-134a refrigerant was used to simulate leaks by drilling small holes into the pipes. The sound of escaping refrigerant was recorded using a microphone, capturing realistic leak audio under experimental conditions. These recordings included both leak events and normal system operation sounds, which were used to help the model learn the acoustic differences between leaking and non-leaking states.

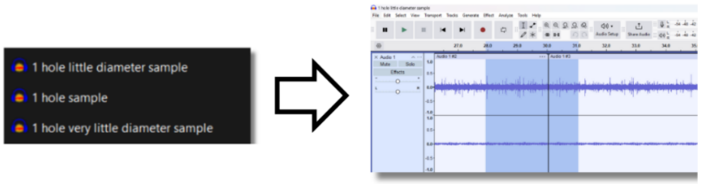

Figure 3.2: Example of Raw Leak Sound Files from Real-World Data

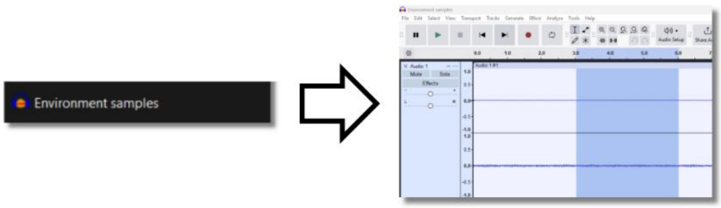

Figure 3.3: Example of Raw Non-Leak Sound (Environment Sound) Files from Real-World Data

By combining generated and real-world audio data, the model was trained with both simulated and authentic leak scenarios, enhancing its generalization and robustness in practical applications.

Real-World Audio Collection

To collect real-world audio data for this study, a controlled experiment was conducted using a refrigerant-based system with R-134a. Simulated leaks were created by drilling small holes of varying sizes into the refrigerant pipes. This setup allowed the capture of realistic leak sounds in a laboratory environment. In addition to leak recordings, environmental sounds were also recorded to represent normal operation without any leakage.

The raw audio consisted of three leak recordings (each about 45 seconds long) and one environmental recording (about 1 minute and 47 seconds). These files were initially saved in .aup3 format. Since these formats were not directly usable for machine learning training, the audio data was preprocessed into short segments ranging from 2 to 5 seconds. Random segments were selected from both the leak and normal recordings and exported as .WAV files for further processing.

The purpose of this data preparation was to create a diverse and balanced dataset, enabling the machine learning model to recognize refrigerant leak patterns under various conditions. By training the model on both leak and non-leak segments, the system can effectively distinguish between normal and abnormal acoustic signals, improving its performance in real-world applications.

Synthetic Leak Signal Generation

To simulate refrigerant leak conditions in the absence of real-world recordings, synthetic leak signals were generated using theoretical calculations. When a refrigerant leaks from a pipe, the escaping flow creates turbulence and pressure disturbances that emit sound waves. These sounds can be characterized by their frequency, which depends on properties such as refrigerant flow velocity, leak size, and fluid dynamics within the pipe.

The process begins with estimating the flow velocity (v) of the refrigerant, calculated using the equation $v = \frac{\dot{m}}{\rho A}$ where $\dot{m}$ is the mass flow rate, ρ is the refrigerant density, and A is the cross-sectional area of the pipe. For example, given a mass flow rate of 5 kg/s, R-134a density of 3.1 kg/m³, and a pipe diameter of 0.079 meters, the calculated flow velocity is approximately 326.53 m/s.

Next, the Reynolds number ($Re = \frac{\rho v D}{\mu}$) is calculated to determine the type of flow in the pipe. With the given values, the resulting Re is approximately 87,181 indicating turbulent flow conditions.

Using the flow velocity, the leak-induced sound frequency (f_{leak}) is then computed using the jet noise equation: $f_{\text{leak}} = S_t \cdot \frac{v}{d}$, where St is the Strouhal number (typically 0.23), and d is the diameter of the leak. For a 5 mm leak, the estimated sound frequency is 15,020.38 Hz. This frequency is within the ultrasonic range and serves as a target signal to generate synthetic leak audio.

To simulate real-world acoustic conditions, the generated pure tone leak signals (e.g., 15 kHz) were mixed with background environmental noises such as air compressor sounds or ambient factory noise. This mixing process ensured that the machine learning model could train on audio samples that closely resemble those encountered in operational environments.

Data Preparation for Model Training

To train the machine learning model effectively, the collected data was carefully organized into three groups. These groups help the model learn to recognize refrigerant leaks and distinguish them from normal sounds. Additionally, a separate dataset was created to test the model after training, ensuring that it can correctly classify new sounds.

Data Organization: Each category (Generate data and Real world data) is further divided into three groups:

1.Generated Data

o Leak Sound Group: 50 audio files, each lasting 2–5 seconds, containing simulated refrigerant leak sounds.

o Not Leak Sound Group: 50 audio files, each lasting 2–5 seconds, containing normal environmental sounds with no leaks.

o Mixed Sound Group: 12 audio files containing a combination of leak and normal sounds for additional model testing.

2.Real-World Data

o Leak Sound Group: 45 audio files, each lasting 2–5 seconds, recorded from actual refrigerant leaks.

o Not Leak Sound Group: 45 audio files, each lasting 2–5 seconds, containing normal environmental sounds with no leaks.

o Mixed Sound Group: 50 audio files, each lasting 2–5 seconds, combining both leak and normal sounds for model testing.

Purpose of Data Preparation

This structured data preparation is important for several reasons:

o Training Accuracy – By separating leak and non-leak sounds, the model can learn to recognize key differences between the two.

o Realistic Testing – The mixed sound group helps evaluate the model's ability to detect leaks in real-world situations.

o Model Generalization – Using both generated and real-world data ensures that the model performs well in various conditions.

This process improves the reliability of the model, making it more effective at detecting refrigerant leaks in different environments.

3.3 Feature Extraction

FFT, STFT, and Wavelet Techniques

Analyzing signals in the frequency domain is crucial for projects involving sound analysis, such as detecting refrigerant leaks. Understanding the frequency characteristics helps identify patterns associated with these leaks. Techniques like Fast Fourier Transform (FFT), Short-Time Fourier Transform (STFT), and Wavelet Transform (WT) are used to break down signals into their frequency components. These methods allow for the extraction of important features necessary for machine learning models to predict leaks accurately. Each method provides different information, and their outputs are combined into a single dataset.

Fast Fourier Transform (FFT)

Theory - The Fast Fourier Transform (FFT) is an efficient algorithm for calculating the Discrete Fourier Transform (DFT). It converts a signal from the time domain (how it changes over time) into the frequency domain (how much of each frequency is present). This transformation is fundamental for identifying key frequency components and analyzing periodicities within a signal. The DFT is mathematically defined by an equation that uses complex exponentials to represent sinusoidal components, converting N time-domain samples into N frequency components. An inverse DFT can reconstruct the original signal.

Application in the Project

The FFT feature extraction process involves several steps:

1. Loading the audio file.
2. Converting the time-domain signal to the frequency domain using FFT.
3. Filtering out frequencies beyond a set limit, such as 25 kHz.
4. Dividing the remaining frequency range into a specific number of bins (n_bins).
5. Calculating the sum of amplitudes within each bin. The result is a set of numerical features representing the frequency information, which helps identify the characteristic frequencies of refrigerant leaks

Short-Time Fourier Transform (STFT)

Theory - The Short-Time Fourier Transform (STFT) builds upon the FFT by analyzing how the frequency content of a signal evolves over time. It achieves this by examining localized sections of the signal, making it especially useful for non-stationary signals where frequencies change. The STFT applies a window function (like Hamming or Hanning) to segments of the signal before performing the Fourier Transform. This provides a time-frequency representation of the signal.

Application in the Project

STFT is applied to capture time-based variations in leak sounds. The process is as follows:

1. Loading the audio file.
2. Computing the STFT, which involves dividing the signal into time segments and applying FFT to each.
3. Removing frequencies above a defined threshold.
4. Dividing the frequency range into n_bins.
5. Calculating the mean amplitude for each bin across time. This yields features that capture both frequency and temporal information.

Wavelet Transform (WT)

Theory - Unlike FFT and STFT, which use fixed frequency bins, the Wavelet Transform (WT) decomposes a signal using wavelets—small, wave-like functions—at multiple resolutions. These wavelets can be scaled and shifted, allowing the WT to capture both short, high-frequency events and long, low-frequency events effectively. The Continuous Wavelet Transform (CWT) defines this process mathematically using scaling and translation parameters. The Discrete Wavelet Transform (DWT) is a simplified version that analyzes the signal at specific scales and translations, making it computationally practical.

Application in the Project

1. The Wavelet Transform helps capture diverse signal characteristics in refrigerant leak sounds. The feature extraction steps are:
2. Loading the audio file.
3. Decomposing the signal using the Discrete Wavelet Transform (DWT).
4. For each resulting wavelet coefficient, computing both the mean amplitude and the standard deviation. This produces statistical features describing the signal's breakdown across different frequency levels.

Implementation Workflow

Feature extraction

Feature extraction is an important step in detecting refrigerant leaks using sound. It helps convert raw audio signals into meaningful data that can be used for machine learning models. In this project, three methods were used to extract features from audio recordings

o Fast Fourier Transform (FFT)

o Short-Time Fourier Transform (STFT)

o Wavelet Transform

Each method analyzes sound in a different way, helping to improve the accuracy of leak detection. The goal of feature extraction is to represent the key characteristics of leak sounds in a structured form, making it easier to differentiate between leak and normal sounds.

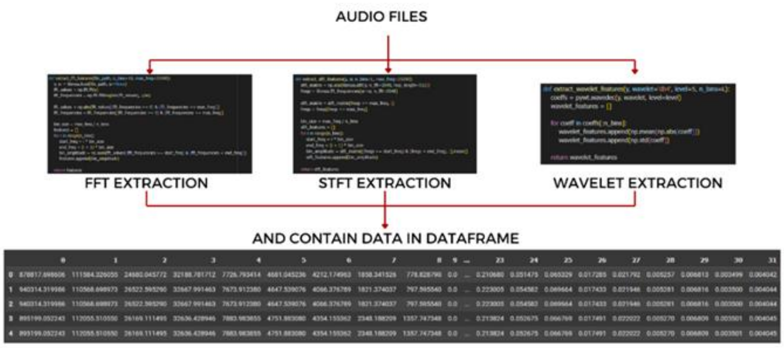

Figure 3.4: Flowchart of Feature Extraction Process for This Model

The feature extraction process can be visualized in a simple flow.

Audio files >> Feature extraction (FFT, STFT, Wavelet) >> Collected features in a dataset

Each method extracts different types of information from the sound, and their outputs are combined into a single dataset for further analysis.

3.4 Model Development

Decision Tree Classifier

Train Model

The machine learning model used in this project is a Decision Tree Classifier. This algorithm works by learning decision rules from the data to classify whether a sound represents a refrigerant leak or not. The decision tree splits the data into smaller groups based on key characteristics, making it easier to identify patterns and classify new sound recordings correctly.

Decision Tree Model Parameters

The Decision Tree model is responsible for classifying sounds as leak or not leak based on extracted features. The following parameters control how the tree is built and how it makes decisions:

1.criterion: Measures the quality of a split ("gini", default).

2.splitter: Determines how the tree chooses where to split the data ("best", default).

3.max_depth: Limits the maximum depth of the tree (None, no limit).

4.min_samples_split: The minimum number of samples required to split a node (2, default).

5.min_samples_leaf: The minimum number of samples required at a leaf node (1, default).

6.max_features: Specifies the number of features to consider when making a split (None, use all features).

These parameters affect how the Decision Tree learns patterns and makes predictions. Keeping most settings at default allows the model to explore all features and create the best decision boundaries.

Data Splitting (Train/Test)

Before training the machine learning model, the dataset is divided into two parts: 80% for training and 20% for testing. This step is crucial because the model needs a large portion of the data to learn patterns, while a separate portion is reserved to evaluate how well the model performs on unseen data. This ensures that the model generalizes well and does not simply memorize the training data.

Model Training and Output

Once the machine learning model was trained using labeled sound data, its performance was evaluated using standard classification metrics. These metrics help determine how accurately the model can identify whether a given sound sample represents a refrigerant leak or not.

The primary metric used was accuracy, which measures the percentage of correct predictions across all test samples. Higher accuracy indicates that the model reliably classifies leak and non-leak sounds.

In addition to accuracy, a detailed classification report was generated, including:

- Precision: The proportion of predicted leak sounds that were actually correct.
- Recall: The proportion of actual leak sounds that were successfully identified.
- F1-score: A balanced metric that combines precision and recall.
- Support: The number of samples for each class in the dataset (leak or not leak).

The final output of the model is a simple classification result for each sound file:

- Leak: The model predicts that the sound indicates a refrigerant leak.
- Not Leak: The model predicts that the sound does not represent a leak.

These outputs provide a practical tool for real-time leak detection, enabling early fault identification and potential system maintenance before larger problems occur.

4. Experimental Results

4.1 Model Performance Overview

Evaluating the performance of the trained model is crucial to understand its accuracy, reliability, and ability to generalize to unseen data. The model was tested using two datasets: Generated Data and Real-World Data. In addition, a separate set of unseen data was used to further assess the model's generalization capability.

The model trained on Generated Data achieved perfect accuracy (100%) during testing, with high precision, recall, and F1-scores. However, when evaluated on unseen data, its accuracy dropped to 66.67%, suggesting possible overfitting — the model performed well on familiar data but struggled to adapt to new, unseen examples.

In contrast, the model trained on Real-World Data initially showed a lower accuracy of 72.22% during testing, but significantly outperformed the Generated Data model on unseen samples. It correctly classified 49 out of 50 unseen files, achieving an accuracy of 98%, along with strong F1-scores and balanced precision-recall values.

These results highlight the importance of training on realistic, diverse audio samples. Although synthetic data is useful in early-stage development, real-world recordings are more effective in producing models that generalize well and perform reliably in practical conditions.

4.2 Evaluation on Test and Unseen Data

To assess the model's performance across different conditions, evaluations were conducted using both test data and unseen data for two distinct training scenarios: Generated Data and Real-World Data. These results provide insight into the model's accuracy and its ability to generalize beyond the data it was trained on.

For the model trained on Generated Data, testing showed excellent results, with 100% accuracy. The classification report demonstrated perfect scores across all key metrics—precision, recall, and F1-score—for both leak and non-leak classes. However, when tested on unseen data (12 new audio files), the model's performance dropped noticeably. It correctly classified 8 out of 12 files, resulting in an accuracy of 66.67%. This gap suggests overfitting, as the model performed well on familiar data but did not generalize effectively to new examples.

In contrast, the model trained on Real-World Data displayed more balanced and practical performance. On the test dataset, it achieved an accuracy of 72.22%, with moderate F1-scores and slightly lower precision for the non-leak class. However, when evaluated on unseen data (50 new files), the model correctly identified 49 out of 50 samples, achieving an impressive 98% accuracy. This indicates strong generalization capability and better real-world applicability compared to the Generated Data model.

Overall, the comparison highlights that while Generated Data is useful for initial training and simulation, Real-World Data significantly enhances the model's ability to perform accurately on new, unseen samples. The results demonstrate that exposure to realistic acoustic environments during training is essential for building robust and deployable machine learning models for refrigerant leak detection.

4.3 Overfitting and Underfitting Validation

In machine learning, overfitting happens when a model learns the training data too well, capturing noise instead of general patterns. This can cause the model to perform perfectly on training data but fail on new, unseen data. On the other hand, underfitting occurs when the model is too simple to capture meaningful patterns, leading to low accuracy in both training and testing. Ensuring a balance between these two is critical for developing a reliable and generalizable refrigerant leak detection model.

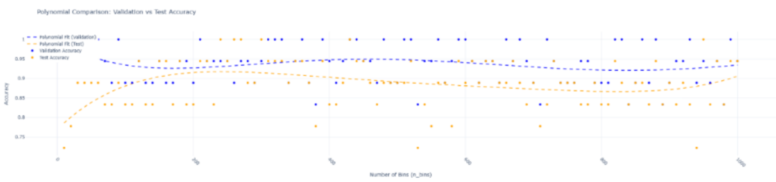

Figure 4.1: Graph Showing Comparison of Accuracy and Validation Accuracy

To address concerns about overfitting and underfitting, the model was validated by determining the optimal number of frequency bins (R) for feature extraction. The validation process involved testing different R values and measuring the difference between validation accuracy and testing accuracy. The results showed that in the range of R = 200 to 240, the difference between validation accuracy and overall model accuracy was the smallest, indicating the best balance between generalization and performance. After further evaluation, R = 220 was selected as the optimal value for the final model, as it provided the most stable and accurate results.

Performance Results

After conducting the validation process, the model was tested using the optimized R value of 220 to ensure balanced performance without overfitting or underfitting. The results show that the model performs accurately and reliably in identifying refrigerant leaks based on sound analysis.

Test Data (Real-World Data)

The model was evaluated using real-world test data, and the results showed an overall accuracy of 94.44%, demonstrating that the model effectively detects leaks in practical conditions. The classification report further confirms the strong performance:

	precision	recall	f1-score	support
0	1.00	0.83	0.91	6
1	0.92	1.00	0.96	12
accuracy			0.94	18
macro avg	0.96	0.92	0.93	18
weighted avg	0.95	0.94	0.94	18

The high precision, recall, and F1-score indicate that the model correctly classifies leak and non-leak sounds with minimal errors. The macro average and weighted average values also show that the model maintains a consistent performance across different categories.

Unseen Data (Real-World Data)

To further evaluate the model's ability to generalize, it was tested on 50 unseen real-world audio files. The model correctly classified all 50 files, achieving 100% accuracy, which confirms its ability to detect refrigerant leaks in new and unknown scenarios.

These results indicate that the validation and tuning process was effective, improving the model's ability to balance learning from the data without overfitting or underfitting. The strong performance on both test data and unseen data suggests that the model is ready for practical applications in refrigerant leak detection, offering a reliable, efficient, and accurate solution for real-world use.

5. Conclusion

5.1 Summary of Findings

This project developed and evaluated a machine learning model for detecting refrigerant leaks using sound data. The model was trained and tested using two types of data: generated audio (simulated leak sounds mixed with environmental noise) and real-world recordings from controlled leak experiments.

In early testing, the model used only basic FFT-based features without dividing the frequency range into smaller segments (bins). This simple approach achieved 100% accuracy on the test data but only 58.33% accuracy on unseen data. This large drop revealed that the model was overfitting—learning the training data too well and failing to generalize to new situations.

To address this, the model was improved by using better feature extraction methods and tuning parameters such as the number of frequency bins (n_bin). After refinement, the updated model trained on generated data still performed well on test data (100% accuracy) but improved to 66.67% on unseen data.

More importantly, when the model was trained using real-world data, it achieved 72% accuracy on the test set and 98% accuracy on unseen data. These results show that real-world data leads to much better generalization, meaning the model can recognize leak sounds it has never seen before.

In summary, the study shows that:

- Simple models may perform well during testing but often fail on unseen data.
- Parameter tuning and advanced feature extraction are key to improving model accuracy.
- Real-world data provides better training examples than synthetic data, leading to stronger model performance in practical scenarios.

5.2 Final Model Performance

To improve model performance, a validation process was conducted to find the optimal number of frequency bins (R) for feature extraction. Different values of R were tested, and results showed that when R was set between 200 and 240, the difference between validation accuracy and testing accuracy was minimized, indicating an improved balance between learning and generalization.

The final model was trained using R = 220, resulting in:

- Accuracy on Testing Data: 89%
- Accuracy on Unseen Data (50 files): 100% (all files correctly classified)

Comparison Before and After Parameter Tuning

Initially, using R = 10, the model achieved:

- Accuracy: 72%
- Unseen Data Accuracy: 98% (49 out of 50 correct)

After tuning the model with R = 220, the accuracy improved by 17%, reaching 89% on test data and 100% on unseen data, confirming significant improvement.

Conclusion

The validation and tuning process significantly enhanced the model's ability to detect refrigerant leaks accurately. The final optimized model achieved higher accuracy and better generalization, ensuring reliable performance in real-world applications. These findings highlight the importance of parameter tuning in machine learning models and confirm that sound-based leak detection is a practical and effective solution for identifying refrigerant leaks in various environments.

5.3 Recommendations for Future Work

In the future, it is important to continue testing the model and fine-tune various parameters to improve its performance. While the model shows promise, optimizing parameters such as feature extraction methods, model settings, and other variables will help enhance its accuracy and reliability. These adjustments will ensure the model performs better in real-world applications, where conditions may vary.

It is important to note that this experiment primarily used data from a water cooler, which represents a small-scale system. However, the ultimate goal of this project is to create a model that can be applied to detect refrigerant leaks in large-scale industrial systems, which are much more complex. In the future, the project should focus on collecting data from real industrial environments, as well as developing the model further to ensure it works efficiently with large systems.

By improving data collection methods and refining the model, it will be better suited for industrial use. The model should also be adapted to handle the challenges specific to large systems, such as various types of equipment, environmental noise, and system variations. This will help reduce the potential issues discussed earlier and make the model more effective in real-world industrial settings.

References

1. Blevins, R. D. (1990). *Flow-induced vibration*. Van Nostrand Reinhold.
2. Dalzochio, J., Kunst, R., Pignaton, E., Binotto, A., Sanyal, S., Favilla, J., & Barbosa, J. (2020). Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges. *Computers in Industry, 123*, 103298. https://doi.org/10.1016/j.compind.2020.103298
3. Kim, M. S., Yun, J. P., & Park, P. (2021). An explainable convolutional neural network for fault diagnosis in linear motion guide. *IEEE Transactions on Industrial Informatics,* 17(6), 4036–4045. https://doi.org/10.1109/TII.2020.3012989
4. Lei, Q., Zhang, C., Shi, J., & Chen, J. (2022). Machine learning based refrigerant leak diagnosis for a vehicle heat pump system. *Applied Thermal Engineering, 215*, 118524. https://doi.org/10.1016/j.applthermaleng.2022.118524
5. Mallat, S. (2008). *A wavelet tour of signal processing: The sparse way* (3rd ed.). Academic Press.
6. Mtibaa, A., Sessa, V., Guerassimoff, G., & Alajarin, S. (2024). Refrigerant leak detection in industrial vapor compression refrigeration systems using machine learning. *International Journal of Refrigeration, 161*, 51–61. https://doi.org/10.1016/j.ijrefrig.2024.02.016
7. Oppenheim, A. V., & Schafer, R. W. (2010). *Discrete-time signal processing* (3rd ed.). Pearson.

8. Rabiner, L., & Schafer, R. (1978). *Digital processing of speech signals*. Prentice Hall.
9. Schlichting, H., & Gersten, K. (2017). *Boundary-layer theory* (9th ed.). Springer.
10. Tam, C. K. W. (1998). Jet noise: Since 1952. *Theoretical and Computational Fluid Dynamics*, 10(1), 393-405. https://doi.org/10.1007/s001620050072
11. White, F. M. (2011). *Fluid mechanics* (7th ed.). McGraw-Hill.