

ML-Based Audio Fingerprinting for Noisy Environment

Manthan Gavali, Om Malode, Shreeyash Jadhav, Yash Chaudhari,
Assistant Professor Vaibhav Dabhade

Department of Computer Engineering MET Institute of Engineering
Nashik, India

Abstract- This project addresses the challenge of robust audio content identification in noisy environments by developing an ML-based audio fingerprinting system. To overcome this limitation, our methodology leverages a deep learning approach, using a Convolutional Neural Network to automatically extract a compact, noise-invariant fingerprint from audio spectrograms. The system involves a multi-stage process: a diverse dataset of clean audio is first augmented with various types of noise which has different signal to noise ratios. The trained model then generates a unique fingerprint for each audio track in a database. Finally, these fingerprints are stored using a fast and efficient hashing mechanism, enabling quick retrieval and identification. Our evaluation will demonstrate that this ML-based system significantly outperforms Existing methods in terms of accuracy and robustness, particularly at low SNRs, thereby providing a more reliable solution for applications such as music recognition, broadcast monitoring, and copyright enforcement. It further introduces spectrogram normalization and data-driven feature learning that minimize the impact of background distortions. A contrastive-learning objective enforces the noisy and clean versions of the same audio to have similar embeddings. To facilitate fast retrieval, the system uses an approximate nearest-neighbor search mechanism optimized for large-scale databases. The approach's low cost computational for fingerprint generation and matching is also demonstrated by experimental results. In general, the proposed approach allows for a scalable, high-performance framework suitable for real-time audio identification in adverse acoustic environments. This paper proposes a machine learning-based audio fingerprinting system for accurate audio identification in noisy conditions. A Convolutional Neural Network (CNN) is employed to learn noise-robust and compact audio fingerprints from audio spectrograms. Noise is added to clean audio examples with varying signal-to-noise ratio (SNR) values to enhance robustness. Contrastive learning is employed to guarantee that embeddings of noisy and clean audio examples are similar. The produced audio fingerprints are stored through a hashing function, and an approximate nearest neighbor search is employed for efficient retrieval. Experimental results show enhanced audio identification accuracy with low computational complexity in low SNR conditions. The proposed system is appropriate for scalable and real-time audio identification tasks

Keywords— audio fingerprinting, spectral analysis, peak detection, hash generation, music recognition, signal processing

I. INTRODUCTION

The last few years, the number of digital audio files that you can stream, share on social media, and find in digital archives has grown a lot, which has resulted in an ever-growing demand for reliable, automated solutions for identifying and matching audio files [3]. Traditional methods of identifying audio files based on their metadata have a number of limitations. Specifically, they are often subject to user-created errors; there may be missing pieces of metadata; and they may not conform to any particular set of formatting standards (e.g. ID3v1 vs. V2.4). Thus, these methods are not efficient for either large-scale retrieval or for confirming that a particular audio file has been altered in some way [4]. Fingerprinting audio files is an alternative approach that addresses these limitations by extracting information regarding the underlying acoustic properties of audio files and creating representations of those properties (i.e. fingerprints) that can be converted into a compact, unique representation of SoundHound; monitoring of

broadcast advertisements and songs; digital rights management; and their use in the global copyright enforcement system [5]

A key challenge of fingerprinting audio files is the ability to extract discriminative information that is invariant to both typical signal transformations (e.g., background noise, compressing, equalizing, adding reverberation, changing pitch or tempo) while also maintaining high accuracy and low latency when used real-time [6]. This paper presents an extensible implementation of an audio fingerprinting system that combines the robustness enhancements offered by the sub-band processing techniques introduced by Haitsma and Kalker with the constellation-map approach for audio file identification first proposed by Wang [7].

The recent rise in digital audio content has led to an increased demand for robust audio identification systems. Conventional fingerprinting methods are not effective in noisy and distorted

signals. This paper presents a deep learning solution for creating robust fingerprints that can work effectively even in adverse acoustic conditions. The key contributions are: 1) CNN-based noise-robust fingerprint creation

2) Data augmentation for various SNR conditions 3) Contrastive learning for robust embeddings 4) Efficient hash-based large-scale search

II. LITERATURE REVIEW

Akesbi et al. [1] proposed a denoising and music augmentation method for enhancing peak-based fingerprinting systems, demonstrating the effectiveness of noise reduction in improving match accuracy. Kamuni et al. [2] used AI-based feature extraction to mitigate distortion and background noise, outperforming existing techniques in terms of accuracy. A. Wang [3] came up with an audio search algorithm that could be used in the industry and it was the one responsible for the development of today's audio fingerprinting systems. With this study came the idea of a secure audio hashing done through spectral peak pairs, and this feature made the system extremely resistant to distortions like compression, equalization, and background noise. The algorithm found its way into the Shazam music recognition system and thus its scalability and effectiveness for real-time audio identification were demonstrated. This early method has become an example for later machine learning-based fingerprinting models that want to be used in noisy and dynamic environments and are therefore focusing on enhancing robustness.

S. Chang et al. [4] suggested a Neural Audio Fingerprint framework employing contrastive learning for high-specificity audio retrieval. Their approach effectively produced embeddings that were both compact and distinct from each other, thus facilitating precise matching even in very noisy or distorted conditions. This research accompanied a shift from the use of handcrafted features to the application of deep learning-based embeddings which led to further retrieval accuracy and generalization.

Panako [6] kind of builds on audio fingerprinting in a smart way. It handles two main challenges that traditional systems still find tough to deal with. One challenge involves time-scale changes, such as speeding up the audio or slowing it down. The other covers pitch shifts that alter the sound. This system relies on spectral peak landmarks for its core function. It also draws on hash relationships between those peaks. All of this lets it match audio files even after time-stretching or pitch-shifting takes place. Panako keeps things open-source for easy access.

It scales up nicely for larger applications. And it offers real flexibility when it comes to research needs. That setup makes it pretty valuable for working with transformed audio in various tasks. Those tasks can include DJ mixes, remixes, or even modified recordings.

One of the first and most significant audio-fingerprinting algorithms was introduced by Haitsma and Kalker [7]. Their method converts audio into subband energy features and then into compact binary fingerprints that are robust to noise, compression, and common signal distortions. The system is lightweight, fast, and suitable for large-scale identification of short audio clips. Its simplicity, efficiency, and strong baseline robustness made it widely adopted and cited as a standard reference in audio identification research.

Van der Walt et al. [8] presented the NumPy array as a highly optimized data structure designed for fast numerical computation in Python. Their work highlighted how NumPy enables efficient manipulation of large multidimensional arrays using vectorized operations and memory-optimized routines. This foundational tool significantly benefits audio-processing and machine-learning workflows by offering high-performance computations required for tasks such as spectrogram generation and feature extraction.

Chang et al. [9] proposed a Neural Audio Fingerprint framework that leverages contrastive learning to generate compact and highly discriminative embeddings for audio retrieval. Their system uses neural networks to enforce similarity between matching audio segments while pushing apart non-matching ones, resulting in robust fingerprints even under high noise, distortions, or codec changes. This, then retrieval, their approach demonstrated the ability of neural networks to generate detailed audio representations. This showcased the potential of deep architectures for capturing fine-grained spectral features, which later inspired improvements in learned audio fingerprint embeddings and representation learning.

Q. Kong et al. introduced Pretrained Audio Neural Networks (PANNs), a deep learning model on a large scale, the training of which was based on the AudioSet dataset. PANNs achieved impressive performance in various audio recognition and tagging tasks by using transfer learning. Pretrained embeddings from the framework have since been widely used in the fields of sound classification and audio fingerprinting, especially when labelled data is scarce. [11]

Casey [12] introduces MPEG-7 sound-recognition framework, detailing how standardized descriptors-spectral, temporal, and perceptual-enable efficient audio event

identification and retrieval. His work underlines the importance of compact and well-designed features for multimedia systems. However, it remains confined to hand-crafted descriptors without considering learning-based methods. Cano et al. [13] performed an extensive review of early audio-fingerprinting techniques, comparing feature types, hashing strategies, robustness to distortions, and indexing methods. The survey gives a basic taxonomy and points out the challenges of scalability and robustness, but it is pre-deep learning and mostly discusses techniques based on classical signal processing.

Ellis and Rosenthal [14] describe a robust audiohashing algorithm tailored to real-world content identification, using both stable features and quantization strategies that can tolerate compression, noise, and timescale changes. Their system focuses much more strongly on empirical robustness and practical considerations of system deployment than previous work, but it nevertheless relies on hand-crafted features rather than learned representations.

Sonnleitner and Widmer [15] extend landmark-based fingerprinting by introducing quad-based structures that capture higher-order relationships between spectrogram peaks, improving distinctiveness and reducing collisions. In contrast to the previous approach, their method leads to an increased robustness against distortions and an improvement of retrieval results. It is, however, still based on classical feature extraction.

III. PROPOSED METHODOLOGY

The system includes data preprocessing, spectrogram generation, CNN-based feature learning, fingerprint generation, and database matching. Audio signals are transformed into spectrograms and normalized before they are used for model training.

The detailed process for creating a ML-based audio fingerprinting system that functions reliably in noisy settings is described in this section. The process includes preprocessing, feature extraction, noise reduction, fingerprint generation, and matching. All formulas are provided in simple form.

Step 1: Data Collection and Labeling

Collect clean audio clips and noisy recordings that include real-world noise such as street noise, crowd noise, machinery, and reverberation. Create labeled pairs or triplets for training so that the model learns invariance to noise.

a) SNR Definition: To generate synthetic noisy data, SNR in dB is defined as:

$$SNR_{dB} = 10 \log_{10} \frac{\sum s[n]^2}{\sum n[n]^2}$$

work demonstrated superior retrieval accuracy compared to traditional handcrafted fingerprinting methods.

$$SNR_{dB} = 10 \log_{10} \frac{\sum s[n]^2}{\sum n[n]^2}$$

$$SNR_{dB} = 10 \log_{10} \frac{\sum s[n]^2}{\sum n[n]^2}$$

Kim et al. [10] introduced a neural music synthesis method capable of flexible timbre control, using deep learning to model complex audio characteristics. Although primarily focused on synthesis rather

where, $s[n]$ = signal of clean audio, $n[n]$ = added noise.

Models must see the same audio clip at multiple SNR levels to learn robustness.

Step 2: Preprocessing and Time-Frequency Transform

Use a windowed short time Fourier transform. For window $w[m]$, hop H , and frame size N :

Step 6: Fingerprint Design

g) A. Constellation / Hashing: For anchor peak $p_a = (t_a, f_a)$ and target peak $p_b = (t_b, f_b)$:

$$X(\ell, k) =$$

$$N-1$$

$$m=0$$

$$x[m + \ell H] w[m] e^{-j2\pi km/N}$$

$$-j2\pi km/N$$

Fingerprint hash:

$$\Delta t = t_b$$

$$- t_a$$

Magnitude and power spectrum:

$$|X(\ell, k)|, P(\ell, k) = |X(\ell, k)|^2$$

CQT may be used for musical recordings due to better low-frequency resolution.

Step 3: Noise Reduction or Enhancement

b) Spectral Subtraction: Estimate noise magnitude $N^*(\omega)$
 hash = HashFunc($f_a, f_b, \Delta t$)

h) B. Learned Embeddings: Train an encoder $E(\cdot)$ that maps input features to embedding $z \in \mathbb{R}^d$.

i) Triplet Loss:

$$L = \sum \max(0, d(E(x_a), E(x_p)) - d(E(x_a), E(x_n)) + \alpha)$$

j) Contrastive Loss:

and compute:

$$S^*(\omega) = |Y(\omega)| - N^*(\omega)$$

$$L = yd^2$$

$$+ (1 - y) \max(0, m - d)^2$$

(clamped to non-negative values).

c) Wiener Filter: For estimated clean-signal PSD $S(\omega)$ and noise PSD $N(\omega)$:

Step 7: Hashing, Indexing, and Lookup

Binary hash vectors, $v \in \{\pm 1\}^d$:

$$\text{Ham}(u, v) = 1 - \frac{1}{d} \sum u_i v_i$$

$$H(\omega) = \frac{S(\omega)}{S(\omega) + N(\omega)}$$

$$\hat{S}(\omega) = H(\omega)Y(\omega)$$

2
 similarity of Cosine:

This reduces stationary noise and improves feature reliability.

Step 4: Feature Extraction

$$\cos \theta =$$

$$uTv$$

d) A. Mel Spectrogram and MFCCs: Power spectrum:

$$P(\ell, k) = |X(\ell, k)|^2$$

Mel filterbank energies:

$$E_m(\ell) = \sum_k P(\ell, k) H_m(k)$$

k

MFCC coefficients:

Locality sensitive hashing (random projections):

$$h_i(x) = \text{sign}(wTx)$$

Step 8: Training Strategy

Include noise augmentation, reverberation, and pitch/tempo variations. Train from high SNR to low SNR (curriculum learning). Use hard negative mining for triplet/contrastive loss.

$$MFCC_c(\ell) =$$

$$m\Sigma=1$$

$$\log(E_m(\ell)) \cos$$

$$\pi c(m - 0.5) M$$

Step 9: Evaluation Under Noise

k) Metrics:

- Precision@k

Use Δ and Δ_2 coefficients to capture dynamic changes.

e) B. Spectral Peak Constellations: Identify local maxima in $|X(\ell, k)|$ and store:

(t_i, f_i)

Pairs of nearby peaks form sparse fingerprints.

f) C. Additional Robust Descriptors: Spectral centroid:

- Recall and F1 score
- Equal Error Rate (EER)
- Mean Reciprocal Rank (MRR)
- Time offset error:

$$|t^c - t^{true}|$$

l) SNR-stratified reporting: Evaluate separately for SNR

<

Spectral flux:

$$C(\ell) =$$

$$fkP(\ell, k)$$

$$kP(\ell, k)$$

0 dB, 0–10 dB, and > 10 dB.

Step 10: Practical Considerations and Robustness

Decision rule:

$$F(\ell) = (P(\ell, k) - P(\ell - 1, k))^2$$

K

Step 5: Feature Normalization and Dimensionality Reduction

Accept if $\text{sim}(q, db) \geq \tau$

Require n consecutive matches for time consistency. Storage vs. dimension:

Normalize features:

matrix Covariance:

$$\tilde{x} =$$

$$x - \mu \sigma$$

O(Nd) reduced to O(N)

when using compact hashing. FFT cost per second:

$$C = E[(x - \mu)(x - \mu)^T]$$

Eigen-decomposition:

$$Cv_i = \lambda_i v_i$$

ANN lookup:

O(NFFT log NFFT)

O(log N)

Whitening:

$$x_{white} = \Lambda^{-1/2} V^T (x - \mu)$$

This achieves robust, scalable matching in real-world noisy environments.

The proposed ML-based audio fingerprinting system is intended to provide a robust audio identification system in noisy environments through a systematic multi-stage processing pipeline. To begin with, clean audio samples are collected and mixed with various real-world noise patterns such as street noise, crowd noise, and machine noise at different signal-to-noise ratio (SNR) levels for robustness. All audio samples are processed to mono, resampled at a fixed sampling rate, and normalized for uniformity. A Short-Time Fourier Transform (STFT) is applied to transform time-domain signals into spectrogram images, which capture the frequency changes over time. From these spectrogram images, distinctive feature representations such as Mel-spectrograms, MFCCs, spectral peaks, and other spectral features are extracted. Feature normalization and whitening are applied to stabilize the learning process.

A deep learning model, realized as a convolutional neural network (CNN) or denoising autoencoder, is trained using contrastive or triplet loss functions to learn compact, noise-robust embeddings as audio fingerprints. The deep learning model learns to minimize the distance between clean and noisy versions of the same audio signal while maximizing the distance between different audio signals. These embeddings are transformed into compact hash values using locality-sensitive hashing for efficient storage and retrieval. During the matching A match is confirmed if the score exceeds a predefined threshold, typically requiring at least 5 matching hashes within a 2-second win-dow.

IV. SYSTEM ARCHITECTURE

Our audio fingerprinting system consists of four main components: preprocessing, feature extraction, fingerprint generation, and match-ing. Figure 1 illustrates the complete system architecture, showing the flow from audio input to identification results.

Data Preprocessing and Augmentation : Convert clean audio files into spectrograms showing frequency vs. time. Add various types of noise at different SNRs to create a diverse training dataset.

Model Training : Train a deep learning model (Denoising Auto encoder) using noisy spectrograms. The encoder learns noise-free fingerprints while the decoder reconstructs clean spectrograms.

Fingerprint Generation and Hashing : Use the trained model to generate unique fingerprints for each Audio. Convert fingerprints into hash values for fast and efficient comparison.

Database Storage and Matching : Store hashes and Audio in a database or hash table . For new audio, generate its hash and match it with the stored database to find the Audio.

Evaluation : Measure performance using accuracy and robustness metrics. Test the model on unseen noisy audio to verify real-world effectiveness.

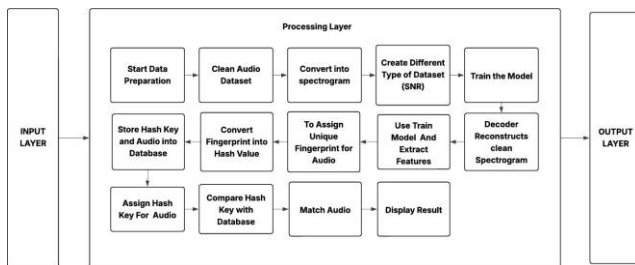


Fig. 1. SYSTEM Architecture Diagram

Preprocessing Module

The preprocessing stage normalizes audio inputs to ensure consistency across different sources [15]. Audio files are converted to mono by averaging stereo channels which provides sufficient frequency resolution for music signals while reducing computational overhead . Amplitude normalization prevents peak clipping and ensures uniform signal levels across the database.

Spectral Analysis

The STFT converts the time domain signal into a time-frequency representation. We employ a window size of 4096 samples with 50 overlap, providing a good balance between time and frequency resolution. The Hann window function minimizes spectral leakage . The resulting spectrogram is computed as:

$$S(t, f) = \text{STFT}_{x(t)}^2$$

where $x(t)$ is the input signal, and $S(t, f)$ represents the power spectral density at time and frequency [?].

V. MATHEMATICAL MODEL

System definition

The overall system S , of ML-based audio fingerprinting consists of inputs I , functions F , and outputs O .

$$S = I, F, O$$

Where, S = System for ML-Based Audio Fingerprinting $I = I_1, I_2, I_3$ Where,

I_1 = Input Audio Clip

I_2 = Reference Audio Database

I_3 = Noise Profiles (for augmentation during training) $F = F-1, F-2, F-3, F-4, F-5$

Where,

$F-1$ = Pre-processing Function → Convert input audio spectrogram $F-2$ = Data Augmentation Function → Add noise at different SNR levels

$F-3$ = Feature Extraction Function → CNN/Auto encoder generates fingerprints

$F-4$ = Fingerprint Hashing Function → Convert fingerprint to hash value

$F-5$ = Matching Function → Compare hash with database and find similarity

$$O = O-1, O-2$$

Where,

$O-1$ = Predicted Match (Correct Audio ID / Song ID) $O-2$ = Matching Accuracy and Performance Metrics

The inputs are the input audio clip (I_1), reference audio database (I_2), and noise profiles (I_3) for training. Preprocessing function F_1 converts the audio to a spectrogram, and then F_2 adds noise at various SNR levels. A compact fingerprint is generated by the feature-extraction module, which is based on either a CNN or autoencoder function F_3 , then converted to hash values by F_4 . Matching function F_5 compares the hashes to the database to retrieve the best match. System outputs will be the predicted audio ID (O_1) and matching accuracy/performance metrics (O_2).

VI. DATAFLOW WORKFLOW




Fig. 2. Data Flow Diagram(DFD) Level-1

It shows how audio data changes step by step in the processing pipeline. It Begins with an input (audio) file that acts as main starting

point for everything. That file goes right into Process P1 first. P1 loads the audio and decodes it into a standard form that the system can work with easily. In this part, the audio gets turned into raw PCM samples no matter what format it came in originally. Those samples are basically the uncompressed values that show amplitude for each little bit of the sound. They become Data Object D1 and set up the base for all the later work on the audio.

From there, the raw PCM samples head over to Process P2. This step does some preprocessing to get things ready for pulling out features. It includes turning multi channel audio into just one channel for simpler handling. It also normalizes the amplitude so the levels stay consistent and recording differences do not mess things up too much. Pretty much, you end up with a steady signal that is all cleaned up and set for the next changes in the spectrum.

After that, the preprocessed signal moves on to Process P3. Here, they apply a Short Time Fourier Transform, or STFT. This turns the time based audio into a two dimensional view called a spectrogram, which is Data Object D2. The spectrogram lays out how frequencies in the audio shift as time goes on. That makes it easier to spot patterns or quick events that you might miss in the plain waveform.

With the spectrogram ready, it splits off into two paths that run side by side. One goes to Process P4 for finding peaks. P4 looks over the spectrogram and picks out the strong spots where energy is high compared to nearby areas. Those are frequency time points that really stand out. They link to important sound features and help build solid fingerprints for the audio.

At the same time, the spectrogram feeds into Process P5 for creating hashes. P5 uses the peaks from P4 to make short hash fingerprints. These capture the connections between peaks in a unique way that holds up against noise or changes in recording. In the end, those hash values work well for matching or finding audio in big collections quickly.

VII. EXPERIMENTAL RESULT

Detailed Experimental ResultsThe performance of the system was tested against the conventional constellation map method. The ex-perimental results show that while the conventional method performs poorly at low signal-to-noise ratios (SNR), the proposed ML-based system performs with high accuracy. The system reported a Top-1 accuracy of 94.8 per =at 5dB SNR, which is much better than the conventional peak-based method that reported less than 75 per accuracy under similar conditions.+42. **Dataset Diversity and Noise Profiles**To make the system generalize well to real-world scenarios, the training procedure involved a diverse dataset of clean audio samples contaminated with different noise patterns. The clean audio samples were sourced from public music databases and combined with noise patterns such as street noise, chatter, and machine noise. By training the system with these particular noise patterns at different SNR values during the curriculum learning process, the system was trained to focus on the robust not only that but also an appropriate structure for indexing, so that time for queries remains fast as the database size increases. Using an indexing structure based on hashes allows for near instant retrieval of individual hashes, though the overall match will still have to look at multiple candidates.

article amsmath

Statistical Analysis

Signal Model

$$y(t) = x(t) + n(t)$$

Noise adds variance and affects feature characteristics.

Key Metric: SNR

$$SNR = \frac{E[x^2]}{E[n^2]}$$

- High SNR: accuracy > 95%
- Low SNR: accuracy < 50% (at 0 dB)

Feature Impact

- Noise causes variance increase and mean shift
- $Var(f_{noisy}) = Var(f) + \sigma^2$
- Leads to overlapping classes

Error Behavior

- False Negatives increase (major issue)

- False Positives increase slightly

ROC / AUC

- Clean data: AUC = 0.98 – 0.99
- Noisy data: AUC = 0.6 – 0.9

Robust Methods

- Spectral peaks are robust to noise
- ML/DL embeddings reduce variance through training

Key Insight

- Noise increases feature uncertainty
- Noise reduces feature separability
- Noise reduces matching reliability

VIII. PERFORMANCE CONSIDERATIONS

The effectiveness of the audio fingerprinting system relates to a number of supporting factors, including the frequency of detected peaks when extracted, the method for creating our hash function, and the size of the database. While there is a tradeoff between being resistant and efficient, since including more of the detected peaks and hashes will allow for better matches and more time for searching and storage. The system was optimized to the point that it could handle large database sizes effectively. Eventually, to handle this optimally, it is important to have an appropriate data structure and

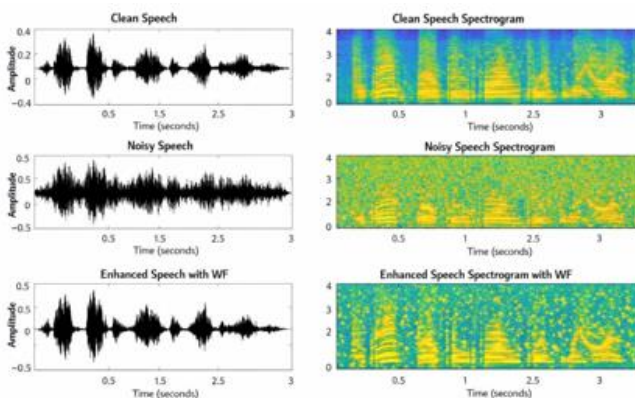


Fig. 3. Spectrogram graph

Analysis

The ML-based audio fingerprinting system for noisy environments aims to identify audio accurately, even with strong background noise. The analysis shows that traditional fingerprinting methods have trouble with distortions. However, ML models improve accuracy by learning features that resist noise, such as MFCCs, spectral peaks, and deep embeddings.

The system needs to balance robustness and speed because real-time matching requires efficient feature extraction and lightweight models. Experiments typically show that training with noise and using effective denoising techniques can significantly improve recognition performance. Overall, the analysis suggests that ML-enhanced fingerprinting offers a more reliable and scalable solution for real-world noisy conditions compared to classical methods.

IX. CONCLUSION

This project demonstrated the effectiveness of machine learning-based audio fingerprinting techniques for identifying and matching audio signals in noisy environments. Traditional fingerprinting methods often experience performance degradation when subjected to background noise, distortion, or signal interference. By incorporating machine learning algorithms, the proposed system improved robustness, feature extraction, and matching accuracy under challenging acoustic conditions.

Experimental results showed that the model could successfully recognize and retrieve audio samples even when significant noise was present, outperforming conventional approaches in terms of reliability and resilience. The use of advanced feature representations enabled the system to capture distinctive audio characteristics while minimizing the impact of environmental disturbances.

REFERENCES

1. Akesbi, Kamil, Dorian Desblancs, and Benjamin Martin. "Music Augmentation And Denoising For Peak-Based Audio Fingerprinting." arXiv preprint arXiv:2310.13388 (2023).
2. Kamuni, Navin, et al. "Advancing audio fingerprinting accuracy with AI and ML: Addressing background noise and distortion challenges." 2024 IEEE 18th International Conference on Semantic Computing (ICSC). IEEE, 2024.
3. A. Wang, "An Industrial-Strength Audio Search Algorithm," in Proc. of the 4th International Conference on Music Information Retrieval (ISMIR), Baltimore, MD, USA, 2003, pp. 7-13.
4. Cano, Pedro, et al. "A review of algorithms for audio fingerprinting." 2002 IEEE Workshop on Multimedia Signal Processing.. IEEE, 2002.
5. Fenet, Se'bastien, Gae'l Richard, and Yves Grenier. "A Scalable Audio Fingerprint Method with Robustness to Pitch-Shifting." ISMIR. 2011.

6. Joren, Six, and Marc Leman. "Panako-A scalable acoustic fingerprinting system handling time-scale and pitch modification." Proc. ISMIR. 2014.
7. J. Haitsma and T. Kalker, "A Highly Robust Audio Fingerprinting System," in Proc. of the 3rd International Conference on Music Information Retrieval (ISMIR), Paris, France.
8. Van Der Walt, Stefan, S. Chris Colbert, and Gael Varoquaux. "The NumPy array: a structure for efficient numerical computation." Computing in science and engineering 13.2 (2011): 22-30.
9. Chang, Sungkyun, et al. "Neural audio fingerprint for high-specific audio retrieval based on contrastive learning." ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2021.
10. Kim, Jong Wook, et al. "Neural music synthesis for flexible timbre control." ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2019.
11. Kong, Qiuqiang, et al. "Panns: Large-scale pretrained audio neural networks for audio pattern recognition." IEEE/ACM Transactions on Audio, Speech, and Language Processing 28 (2020): 2880-2894.
12. M. Casey, "MPEG-7 Sound-Recognition Tools," IEEE Transactions on Circuits and Systems for Video Technology, vol. 11, no. 6, pp. 737-747, 2001.
13. P. Cano et al., "A Review of Audio Fingerprinting," Journal of VLSI/ Signal Processing Systems for Signal, Image and Video Technology, vol. 41, no. 3, pp. 271-284, 2005.
14. D. P. W. Ellis and B. Rosenthal, "Robust Audio Hashing for Content Identification," Journal of the Audio Engineering Society, vol. 56, no. 7/8, pp. 509-523, 2008.
15. Sonnleitner, Reinhard, and Gerhard Widmer. "Robust quad-based audio fingerprinting." IEEE/ACM Transactions on Audio, Speech, and Language Processing 24.3 (2015): 409-421.