

Segment Transformer: AI-Generated Music Detection via Music Structural Analysis

Yumin Kim* and Seonghyeon Go*

MIPPIA Inc.

{kym, gsh}@mippia.com

Abstract—Audio and music generation systems have been remarkably developed in the music information retrieval (MIR) research field. The advancement of these technologies raises copyright concerns, as ownership and authorship of AI-generated music (AIGM) remain unclear. Also, it can be difficult to determine whether a piece was generated by AI or composed by humans clearly. To address these challenges, we aim to improve the accuracy of AIGM detection by analyzing the structural patterns of music segments. Specifically, to extract musical features from short audio clips, we integrated various pre-trained models, including self-supervised learning (SSL) models or an audio effect encoder, each within our suggested transformer-based framework. Furthermore, for long audio, we developed a segment transformer that divides music into segments and learns inter-segment relationships. We used the FakeMusicCaps and SONICS datasets, achieving high accuracy in both the short-audio and full-audio detection experiments. These findings suggest that integrating segment-level musical features into long-range temporal analysis can effectively enhance both the performance and robustness of AIGM detection systems.

I. INTRODUCTION

Recent advances in generative models have enabled high-quality content synthesis across diverse modalities, including images [1], video [2], audio [3], [4], and natural language [5], [6]. In particular, the emergence of intuitive and user-friendly tools such as Suno¹ and Udio² has made AI-based audio and music generation highly accessible, significantly lowering the barrier for the creation of AI-generated music (AIGM). This growing trend has underscored the urgent need for technical solutions capable of reliably distinguishing AIGM from human-composed music and other manipulated audio content, to prevent unethical use and protect music intellectual property [7].

However, existing AI-generated audio detection [8]–[12] methods are often limited in local scope due to CNN architecture, making it hard to extract structural dependencies in full music. There's the SpecTTTra model suggested in [13], which successfully proposed a method to address this issue and extract long-range dependencies from the entire audio. However, it shows limited utilization of musical information such as BPM and downbeat for analyzing musical structure.

To address these challenges, we propose a two-stage AIGM detection framework. In the stage-1, which detects AIGM from short audio segments, we begin by extracting meaningful

representations from each segment. Specifically, we pair our Cross-Attention-based Transformer decoder (AudioCAT) with several self-supervised audio encoders [14]–[17] to implement a specialized detection approach for each feature extractor. We further integrate a pre-trained audio-effect encoder such as FXencoder [18] into our Transformer-based architecture (FXencoder-Segment).

While various methodologies [19] exist for extracting high-quality embeddings from short audio slices, the model's strategy of directly mapping these slices to Transformer token embeddings still has potential for improvement. In stage 2, we slice full audio using beat tracking and apply our Segment Transformer model. It is designed to capture global song features through structural analysis. This approach improves the model's ability to understand relationships between musical features, enabling effective differentiation between the intentional compositional patterns of human musicians and the AI. Additionally, the segment-based approach maintains awareness of musical form while efficiently processing complete compositions, allowing the detection system to identify inconsistencies in musical structure development and motif progression throughout the entire music. Our main contributions are as follows:

(1) Segment Detection Architecture: We propose AudioCAT, a cross-attention-based Transformer decoder architecture that employs feature extractors as encoders to derive meaningful representations from short audio segments. Additionally, we adopt pre-trained FXencoder, which encodes only audio-effect information, to demonstrate the potential of diverse feature extractors for effective AIGM detection.

(2) Segment Transformer Architecture: We propose a Segment Transformer model that employs beat-aware audio segmentation mechanisms to preserve temporal relationships and capture global song features through structural analysis of complete compositions.

(3) Novel Stage-2 Framework: We employ two segment detection models that extract audio features from short segments, while the Segment Transformer captures structural features from longer pieces. By combining these approaches for comprehensive musical analysis, we achieve superior detection performance and establish a new benchmark for full-audio AIGM detection.

*Equal contribution.

¹<https://suno.com/>

²<https://udio.com/>

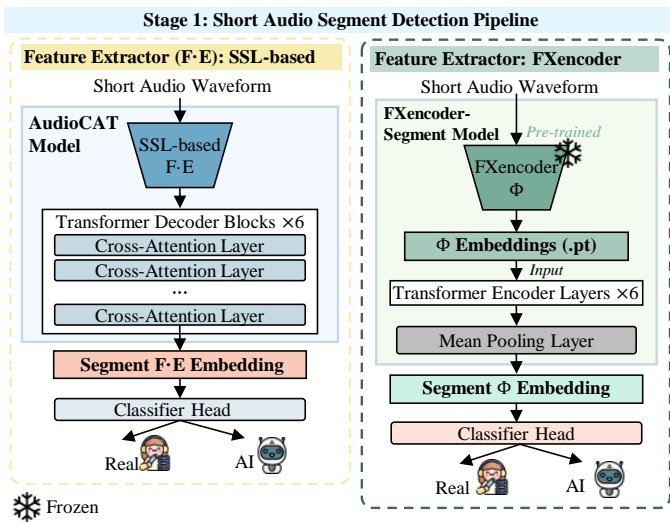


Fig. 1. Overview of the proposed Stage-1: Short audio segment detection pipeline. We employ two types of feature extractors: self-supervised learning (SSL)-based models (Left) and an audio-effect-focused FXencoder [18](Right).

II. METHODS

A. Stage 1: Short Audio Segment Detection Model

As shown in stage 1 of Figure 1, we employ specialized feature extractors as the model’s encoder to detect short audio segments. Specifically, we use self-supervised learning (SSL) models and a pre-trained FXencoder [18] that captures audio-effect information.

1) *AudioCAT*: This architecture integrates an SSL-based feature extraction encoder with a **Cross-Attention**-based Transformer decoder. The feature extraction encoder captures localized frequency patterns and textural elements from audio segments, producing spatially confined feature representations. Moreover, this encoder can be instantiated with a general-purpose SSL model trained on large-scale audio datasets or replaced by a task-specific encoder. In our architecture, the decoder latent states are used as queries, while the encoder feature maps serve as keys and values, allowing the decoder to align its internal contextual representation with localized audio features. By leveraging the cross-attention mechanism to learn sequential representations, the model enhances its discriminative capacity, after which a classification head performs the AI-generated music (AIGM) detection task.

Why Apply Cross-attention. Cross-attention allows the decoder to strategically integrate features with its own internal representations, achieving better integration of local textural and global contextual information, where even short music segments contain crucial global patterns that are essential for effective AIGM detection performance.

2) *FXencoder-Segment Model*: We hypothesize that extracting audio-effect (FX) and music-specific features provides greater advantages for AIGM detection than relying on general audio representations. Consequently, we employ the FXencoder [18]—a specialized model originally designed for

other audio tasks within our self-attention based Transformer encoder architecture. We use the pre-trained FXencoder to extract embedding vectors as input, keeping the feature extractor frozen during training so its representations remain fixed. While the cross-attention mechanism in AudioCAT attempts to “strategically combine” these features with decoder representations, this approach can be inefficient with a fixed FXencoder because the query remains. Therefore, a pure self-attention mechanism based on Transformer encoder layers is more appropriate, allowing all frozen features to directly interact and thereby effectively uncover relationships and patterns within the pre-trained embedding space.

B. Feature Extractor for Segment Detection Model

As shown in Figure 1, we selected both self-supervised learning (SSL) based pre-trained models and the FXencoder [18], which we determined to provide essential characteristics for AIGM detection.

The AudioCAT architecture is compatible with any feature extractor that produces meaningful embeddings from short audio segments. In our methods, we use Wav2vec [20], Music2vec [21], and MERT [19] as the feature extraction encoders. The feature extractor models used in our experiments are described as follows.

Wav2vec 2.0. Existing spoofing detection [15], [22] models primarily use features such as Mel-spectrogram or Mel-Frequency Cepstral Coefficient (MFCC) [16], while Wav2vec-2.0 [20], [23] processes raw waveforms directly, allowing for the learning of more robust representations. Compared to conventional spoofing detection models [8]–[12], it offers better generalization across diverse speech data. We employ a pre-trained Wav2vec 2.0-Base model that extract 512-dimensional SSL feature representations.

Music2vec. This pre-trained model [21] adopts the pre-training paradigm of the speech version of data2vec-1.0 [24]. Specifically adapted for the music domain, this continuous prediction model employs a 1D Convolutional Embedder to provide waveform input to both the teacher and student Transformer models, enabling more effective representation learning for music data. We use a pre-trained Music2Vec-v1 (95M) model with 95M parameters that extract 512-dimensional feature representations.

MERT. A large-scale SSL-based music understanding model [19] is designed to address the challenges in modeling musical knowledge, particularly tonal characteristics and pitched attributes, which have hindered the extensive exploration of SSL applications in music audio. We use the MERT-v1 (95M) model, which is designed to generalize better across more tasks and has been trained on a dataset to achieve superior performance compared to previous models that extracting 768-dimensional feature representations.

FXencoder. Music contains unique structural and harmonic patterns that fundamentally differ from general audio signals. The FXencoder is designed specifically to extract mixing and mastering features from audio files. Unlike SSL-based models that learn general audio patterns, the FXencoder focuses on

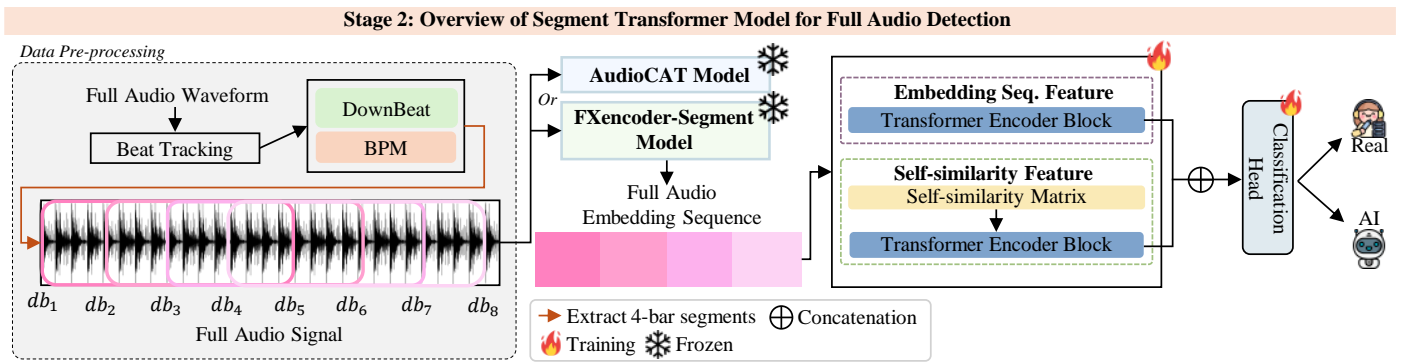


Fig. 2. Overview of the proposed Stage-2: Segment Transformer for full audio segment detection pipeline.

uncovering production details and processing effects in music recordings, adding unique mixing and mastering-related features that are advantageous for AIGM detection. This capacity to extract production-related characteristics is valuable for analyzing how audio was produced and helps distinguish human-produced music from AI-generated compositions. FXencoder extracts a 2048-dimensional feature representation.

The two proposed frameworks are not limited to the specific models presented here; this study can accommodate any architecture capable of extracting meaningful embeddings from short audio segments, ensuring flexibility for future research and alternative feature-extraction approaches.

C. Stage 2: Full-Audio Segment Detection Model

Since real music tracks have varying lengths and diverse structures, full-audio analysis is necessary for robust AIGM detection performance. To overcome these limitations, we introduce the Segment Transformer. We preprocessed audio by dividing tracks into manageable segments with beat tracking model. And we trained structural relationships by two different transformer encoder block. This segmented approach allows for more consistent and accurate detection across different track lengths and musical structures, leading to improved full-audio analysis performance.

Music Segment Preparation. As shown in the stage-2 architecture in Figure 2, our full-audio detection approach begins with processing full audio tracks into musically meaningful segments. We first detect downbeats in the music using a beat-tracking algorithm, which allows us to identify the beginning of each bar. We then split the audio into 4-bar units, creating a sequence of music segments that preserves the natural rhythmic structure of the composition. These musically informed segments ensure that our model processes the audio in a way that respects its musical organization.

Segment Transformer. For full-audio analysis, we introduce the Segment Transformer, which processes sequences of music segments to analyze entire compositions comprehensively. As part of our training approach, while the Segment Transformer is compatible with various feature extractors, we use our segment detection models to capitalize on their domain-specific representations. Their outputs serve as rich segment-

level representations, allowing the Segment Transformer to learn temporal dependencies and structural coherence across complete musical works.

The Segment Transformer employs a dual-pathway architecture with two parallel transformer encoder structures that capture complementary aspects of musical structure. The first encoder processes content embeddings directly, capturing the semantic and acoustic properties of individual music segments. The second encoder analyzes global structural patterns with a self-similarity matrix between all segment embeddings. This method is used in many musical structure analysis [25], enabling the model to understand repetitive structures, variations, and overall compositional organization. The outputs from both encoders are then concatenated to form a unified representation that combines local segment characteristics with global structural information. This provides a comprehensive understanding of the entire musical composition for more accurate AIGM detection.

III. EXPERIMENTS AND RESULTS

A. Datasets

We used the FakeMusicCaps [26] and SONICS [13] datasets for training and evaluating the stage-1 segment detection model. The FakeMusicCaps dataset consists of 10-second short audio generated by five different Text-to-Music (TTM) models, each employing distinct audio compression techniques and text conditioning mechanisms to produce diverse musical outputs from the same textual descriptions. Specifically, MusicGen [35], MusicLDM [36], AudioLDM2 [4], Stable Audio Open [37], and Mustango [38]. The dataset contains a total of 5,373 real and 27,605 AI-generated tracks. The SONICS dataset comprises 48,090 real tracks collected dynamically from YouTube using lyrics, titles, and artist metadata sampled from the Genius Song Lyrics dataset, and 49,074 AI-generated tracks, produced with SunoSuno¹ and Udio² generative models. All tracks have an average duration of 176 seconds. Moreover, for model training, both real and AI-generated tracks were resampled to 16 kHz—consistent with the baseline experiments [13]—and the dataset was configured and used in the same way.

TABLE I

PERFORMANCE COMPARISON OF OUR FEATURE EXTRACTOR-BASED MODELS AND CONVENTIONAL AI-GENERATED MUSIC DETECTION APPROACHES IN THE STAGE 1 PIPELINE: ON FAKEMUSICCAPS (LEFT) USING 10-SECOND AUDIO CLIPS, AND ON SONICS (RIGHT) USING 5-SECOND SEGMENTS FROM EACH TRACK. BEST: **BOLD**; SECOND BEST: UNDERLINE.

FakeMusicCaps Dataset [26]							SONICS Dataset [13] (5-sec Segment)						
<i>Baselines</i>							<i>Baselines</i>						
Model	ACC	Prec.	Recall	F1	AUC	Spec.	Model	ACC	Prec.	Recall	F1	AUC	Spec.
ResNet18 [27]	0.924	–	–	0.924	–	–	ConvNeXt [28]	–	–	0.82	0.90	–	0.98
SeNet [29]	0.923	–	–	0.923	–	–	ViT [30]	–	–	0.80	0.79	–	0.79
MobileNet [31]	0.968	–	–	0.968	–	–	EfficientViT [32]	–	–	0.78	0.87	–	0.98
CNN+LSTM [33]	0.916	–	–	0.917	–	–	SpecTTTra- γ [13]	–	–	0.63	0.76	–	0.98
Q-SVM [34]	0.914	–	–	0.924	–	–	SpecTTTra- β [13]	–	–	0.69	0.78	–	0.94
ViT [30]	0.854	–	–	0.852	–	–	SpecTTTra- α [13]	–	–	0.71	0.80	–	0.92
<i>(Ours) Stage 1: Short Audio Segment Detection</i>							<i>(Ours) Stage 1: Short Audio Segment Detection</i>						
Feature Extractor	ACC	Prec.	Recall	F1	AUC	Spec.	Feature Extractor	ACC	Prec.	Recall	F1	AUC	Spec.
Wav2vec 2.0	<u>0.980</u>	<u>0.988</u>	<u>0.988</u>	<u>0.988</u>	<u>0.944</u>	<u>0.939</u>	Wav2vec 2.0	0.995	0.991	0.999	<u>0.995</u>	<u>0.996</u>	0.990
Music2vec	0.919	0.948	0.956	0.952	0.937	0.730	Music2vec	<u>0.996</u>	<u>0.997</u>	<u>0.997</u>	0.997	1.000	0.996
MERT [19]	0.994	0.996	0.997	0.996	0.998	0.980	MERT	0.997	0.998	0.996	0.997	1.000	0.998
FXencoder	0.860	0.948	0.881	0.913	0.924	0.751	FXencoder	0.971	0.950	0.984	0.966	0.993	0.961

B. Implementation Details

We split the FakeMusicCaps dataset into training, validation, and test sets using an 8:1:1 ratio. For data augmentation, we randomly applied pitch shifting (± 2 semitones) and time stretching (0.8 \times and 1.25 \times) with a 50 % probability was performed exclusively during the training of the SSL-based feature extractor in the segment detection model. Because FXencoder is pre-trained to extract audio features and used with frozen weights, applying augmentations such as pitch shifting or time stretching could distort the original audio patterns. Therefore, we applied data augmentation only to the SSL-based feature extractors, whose weights are updated during training. For our segment detection experiments, input audio was resampled to 16 kHz for Wav2vec and Music2vec, 24 kHz for MERT, and 44.1 kHz stereo for the FXencoder. AudioCAT was trained for 30 epochs with a batch size of 8 using binary cross-entropy loss, while the FXencoder was trained for 50 epochs with a batch size of 32 using focal loss. All models were optimized with the Adam optimizer (learning rate 1×10^{-5} , weight decay 1×10^{-6}) and employed early stopping.

For the Segment Transformer model, we used the SONICS dataset. Each track was divided into a variable number of segments based on its structure, and these segments were then padded or cropped to a fixed sequence length of 48 for input to the transformer model. To extract downbeat information, we used the Beat this! [39] model. The downbeat sequence was quantized into an arithmetic sequence aligned with the music’s structural rhythm. This model was trained for 50 epochs using binary cross-entropy loss and the Fused Adam optimizer (learning rate of 1×10^{-5} , a weight decay of 1×10^{-6}), and applied early stopping.

All models were trained on a single NVIDIA RTX 5090 GPU. We evaluated their performance on the test set using the following metrics: Accuracy, Precision (Prec.), Recall

(Sensitivity), F1-score, area under the ROC curve (AUC), and Specificity (Spec.).

C. Segment Detection Performances

As shown in Table I, the models proposed in this study outperform the CNN-based MobileNet [31], which represents the state of the art (SOTA) on the FakeMusicCaps dataset, in every evaluation metric. Our four proposed models consistently outperform the baseline on nearly every metric. FXencoder [18], which emphasizes mixing and mastering characteristics, shows relatively lower performance compared to other models, it still outperforms the ViT [30] baseline in both accuracy and F1 score. Despite these mixed results, FXencoder demonstrates certain music structural advantages in full-audio detection experiments. On the SONICS dataset, our approaches again demonstrate uniformly superior performance. In particular, MERT [19] achieves optimum results among the models, owing to its rich representations through large-scale pre-training on diverse music and audio tasks.

The improvements are achieved by employing domain-specific encoders—SSL models for general audio understanding and FXencoder for music production details—together with well-matched architectural designs.

D. Full Audio Detection Performances

For the full audio detection model Segment Transformer, we trained and evaluated our model using SONICS dataset. We used the benchmark reported in [13], specifically for the task of full audio (120-second) AIGM detection in the SONICS dataset.

These results can be interpreted through the architectural design of our Segment Transformer, which focuses on modeling temporal connections between musical segments. The performance disparity across different feature extractors reveals the importance of music-specific representations for segment-level modeling.

TABLE II

PERFORMANCE COMPARISON OF FULL-AUDIO SEGMENT DETECTION OF SEGMENT-LEVEL DETECTORS ON SONICS USING 120-SECOND SEGMENTS FROM EACH TRACK. BEST: **BOLD**; SECOND BEST: UNDERLINE.

SONICS Dataset [13](120-sec Segment)						
<i>Baselines</i>						
Model	ACC	Prec.	Recall	F1-Score	AUC	Spec.
ConvNeXt [28]	–	–	0.95	0.96	–	0.98
ViT [30]	–	–	0.82	0.89	–	0.98
EfficientViT [32]	–	–	0.94	0.95	–	0.97
SpecTTTra- γ [13]	–	–	0.79	0.88	–	<u>0.99</u>
SpecTTTra- β [13]	–	–	0.86	0.92	–	<u>0.99</u>
SpecTTTra- α [13]	–	–	0.96	0.97	–	<u>0.99</u>
<i>(Ours) Stage 2: Segment Transformer for Full Audio Detection</i>						
Feature Extractor	ACC	Prec.	Recall	F1-Score	AUC	Spec.
Wav2vec 2.0	0.7473	0.7127	0.8726	0.7846	0.7412	0.6075
Music2vec	0.9555	0.9609	0.9546	0.9577	0.9907	0.9566
MERT	0.9992	0.9988	0.9996	0.9992	0.9999	0.9987
FXencoder	<u>0.9955</u>	<u>0.9953</u>	<u>0.9961</u>	<u>0.9957</u>	<u>0.9995</u>	<u>0.9948</u>

Wav2vec 2.0, being primarily designed for speech representations, resulting in significantly degraded performance. In contrast, music-specific feature extractors such as MERT and FXencoder demonstrate exceptional performance, with MERT achieving near-perfect results across all metrics. This suggests that when combined with musically-informed representations, Segment Transformer can effectively learn temporal segment relationships to achieve superior performance.

IV. CONCLUSIONS

In this paper, we presented a novel approach for detecting AI-generated music (AIGM), addressing the needs for music information retrieval domain. Our key contributions are: (1) AudioCAT, a framework that applies cross-attention mechanisms with diverse SSL-based feature extractors and is designed to accommodate any encoder. (2) Full-audio analysis approach using a segment transformer architecture that preserves structural relationships, and a two-stage robust model framework for AIGM detection with strong performance. By considering both complete musical compositions and isolated segments, our approach effectively captures the music structural difference for distinguishing human-composed from AIGM. Experimental results show that our two-stage processing architecture significantly outperforms existing methods. These findings suggest that SSL models designed for the unique characteristics of musical data can serve as an effective means of enhancing detection performance.

Future work could explore alternative methodological approaches, including end-to-end architectures that directly process full-length audio, as well as investigating different fusion strategies for combining segment-level and track-level information. Additionally, exploring other neural architectures beyond transformers may yield further improvements in detection accuracy and computational efficiency. This research marks a meaningful step toward establishing effective detection methods that can keep pace with the rapid growth of generative AI in music, addressing increasing concerns around copyright

and creative authenticity.

REFERENCES

- [1] N. Yu, L. S. Davis, and M. Fritz, “Attributing fake images to gans: Learning and analyzing gan fingerprints,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 7556–7566.
- [2] U. Singer, A. Polyak, T. Hayes, *et al.*, “Make-a-video: Text-to-video generation without text-video data,” *arXiv preprint arXiv:2209.14792*, 2022.
- [3] F. Kreuk, G. Synnaeve, A. Polyak, *et al.*, “Audio-gen: Textually guided audio generation,” *arXiv preprint arXiv:2209.15352*, 2022.
- [4] H. Liu, Y. Yuan, X. Liu, *et al.*, “Audioldm 2: Learning holistic audio generation with self-supervised pretraining,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [5] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, *et al.*, “Improving language understanding by generative pre-training,” 2018.
- [6] J. Li, T. Tang, W. X. Zhao, J.-Y. Nie, and J.-R. Wen, “Pre-trained language models for text generation: A survey,” *ACM Computing Surveys*, vol. 56, no. 9, pp. 1–39, 2024.
- [7] D. G. Widder and M. Hicks, “Watching the generative ai hype bubble deflate,” *arXiv preprint arXiv:2408.08778*, 2024.
- [8] J. Yi, C. Wang, J. Tao, X. Zhang, C. Y. Zhang, and Y. Zhao, “Audio deepfake detection: A survey,” *arXiv preprint arXiv:2308.14970*, 2023.
- [9] X. Liu, X. Wang, M. Sahidullah, *et al.*, “Asvspoof 2021: Towards spoofed and deepfake speech detection in the wild,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 2507–2522, 2023.
- [10] J. Yi, J. Tao, R. Fu, *et al.*, “Add 2023: The second audio deepfake detection challenge,” *arXiv preprint arXiv:2305.13774*, 2023.
- [11] M. Masood, M. Nawaz, K. M. Malik, A. Javed, A. Irtaza, and H. Malik, “Deepfakes generation and detection: State-of-the-art, open challenges, countermeasures, and way forward,” *Applied intelligence*, vol. 53, no. 4, pp. 3974–4026, 2023.
- [12] Y. Patel, S. Tanwar, R. Gupta, *et al.*, “Deepfake generation and detection: Case study and challenges,” *IEEE Access*, vol. 11, pp. 143 296–143 323, 2023.
- [13] M. A. Rahman, Z. I. A. Hakim, N. H. Sarker, B. Paul, and S. A. Fattah, “Sonics: Synthetic or not—identifying counterfeit songs,” *arXiv preprint arXiv:2408.14080*, 2024.
- [14] Y. Ma, R. Yuan, Y. Li, *et al.*, “On the effectiveness of speech self-supervised learning for music,” *arXiv preprint arXiv:2307.05161*, 2023.
- [15] Y. Xie, Z. Zhang, and Y. Yang, “Siamese network with wav2vec feature for spoofing speech detection,” in *Interspeech*, 2021, pp. 4269–4273.

- [16] H. Tak, M. Todisco, X. Wang, J.-w. Jung, J. Yamagishi, and N. Evans, “Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation,” *arXiv preprint arXiv:2202.12233*, 2022.
- [17] R. Castellon, C. Donahue, and P. Liang, “Codified audio language modeling learns useful representations for music information retrieval,” *arXiv preprint arXiv:2107.05677*, 2021.
- [18] J. Koo, M. A. Martínez-Ramírez, W.-H. Liao, S. Uhlich, K. Lee, and Y. Mitsufuji, “Music mixing style transfer: A contrastive learning approach to disentangle audio effects,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2023, pp. 1–5.
- [19] Y. Li, R. Yuan, G. Zhang, *et al.*, “Mert: Acoustic music understanding model with large-scale self-supervised training,” 2024.
- [20] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “Wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [21] Y. Li, R. Yuan, G. Zhang, *et al.*, “Map-music2vec: A simple and effective baseline for self-supervised music audio representation learning,”
- [22] X. Wang and J. Yamagishi, “Investigating self-supervised front ends for speech spoofing countermeasures,” *arXiv preprint arXiv:2111.07725*, 2021.
- [23] Q. Xu, A. Baevski, T. Likhomanenko, *et al.*, “Self-training and pre-training are complementary for speech recognition,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2021, pp. 3030–3034.
- [24] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, “Data2vec: A general framework for self-supervised learning in speech, vision and language,” in *International conference on machine learning*, PMLR, 2022, pp. 1298–1312.
- [25] Y. Shiu, H. Jeong, and C.-C. J. Kuo, “Similarity matrix processing for music structure analysis,” in *Proceedings of the 1st ACM workshop on Audio and music computing multimedia*, 2006, pp. 69–76.
- [26] L. Comanducci, P. Bestagini, and S. Tubaro, “Fakemusicaps: A dataset for detection and attribution of synthetic music generated via text-to-music models,” *arXiv preprint arXiv:2409.10684*, 2024.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [28] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, “A convnet for the 2020s,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 11 976–11 986.
- [29] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [30] A. Dosovitskiy, L. Beyer, A. Kolesnikov, *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.
- [31] A. G. Howard, M. Zhu, B. Chen, *et al.*, “Mobilenets: Efficient convolutional neural networks for mobile vision applications,” *arXiv preprint arXiv:1704.04861*, 2017.
- [32] H. Cai, J. Li, M. Hu, C. Gan, and S. Han, “Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 17 302–17 313.
- [33] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [34] A. K. Singh and P. Singh, “Detection of ai-synthesized speech using cepstral & bispectral statistics,” in *2021 IEEE 4th International Conference on Multimedia Information Processing and Retrieval (MIPR)*, IEEE, 2021, pp. 412–417.
- [35] J. Copet, F. Kreuk, I. Gat, *et al.*, “Simple and controllable music generation,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 47 704–47 720, 2023.
- [36] K. Chen, Y. Wu, H. Liu, M. Nezhurina, T. Berg-Kirkpatrick, and S. Dubnov, “Musicldm: Enhancing novelty in text-to-music generation using beat-synchronous mixup strategies,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2024, pp. 1206–1210.
- [37] Z. Evans, J. D. Parker, C. Carr, Z. Zukowski, J. Taylor, and J. Pons, “Stable audio open,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2025, pp. 1–5.
- [38] J. Melechovsky, Z. Guo, D. Ghosal, N. Majumder, D. Herremans, and S. Poria, “Mustango: Toward controllable text-to-music generation,” *arXiv preprint arXiv:2311.08355*, 2023.
- [39] F. Foscari, J. Schlüter, and G. Widmer, “Beat this! accurate beat tracking without dbn postprocessing,” *arXiv preprint arXiv:2407.21658*, 2024.