

# TOWARDS ASSESSING DATA REPLICATION IN MUSIC GENERATION WITH MUSIC SIMILARITY METRICS ON RAW AUDIO

Roser Batlle-Roca<sup>1</sup> Wei-Hsiang Liao<sup>2</sup> Xavier Serra<sup>1</sup>  
Yuki Mitsufuji<sup>3</sup> Emilia Gómez<sup>1,4</sup>

<sup>1</sup> Music Technology Group, Universitat Pompeu Fabra, Spain <sup>2</sup> Sony AI, Japan <sup>3</sup> Sony AI, USA

<sup>4</sup> Joint Research Centre, European Commission, Spain

roser.batlle@upf.edu

## ABSTRACT

Recent advancements in music generation are raising multiple concerns about the implications of AI in creative music processes, current business models and impacts related to intellectual property management. A relevant discussion and related technical challenge is the potential replication and plagiarism of the training set in AI-generated music, which could lead to misuse of data and intellectual property rights violations. To tackle this issue, we present the Music Replication Assessment (MiRA) tool: a model-independent open evaluation method based on diverse audio music similarity metrics to assess data replication. We evaluate the ability of five metrics to identify exact replication by conducting a controlled replication experiment in different music genres using synthetic samples. Our results show that the proposed methodology can estimate exact data replication with a proportion higher than 10%. By introducing the MiRA tool, we intend to encourage the open evaluation of music-generative models by researchers, developers, and users concerning data replication, highlighting the importance of the ethical, social, legal, and economic consequences. Code and examples are available for reproducibility purposes.<sup>1</sup>

## 1. INTRODUCTION

Significant advancements in generative algorithms for digital art creation are challenging the role of artificial intelligence (AI) in artistic practices. Regarding generative AI in the music domain, there is an increasing discussion related to the use of computational tools in music creative processes [1], the effects on artists' work, existing listening experiences and business models, and the impacts on intellectual property (IP) management [2,3]. A key point is the potential replication and plagiarism of the training set in AI-generated music [3,4], which can lead to data misuse and IP violations.

<sup>1</sup><https://github.com/roserbatlleroca/mira>

The inherent opaque nature of music generation models complicates tracing replications or references to training set samples in AI-generated music, limiting the interpretation of whether generated samples contain replicated fragments. In addition, diffusion models, one of the most popular generative AI architectures, tend to memorise and replicate training data [5–7]. Understanding the behaviour of these models has become critical to address legal issues [8], especially when dealing with data protected by IP rights. This is significant in the music domain as the vast majority of music is protected by authorship and copyright.

Despite multiple claims emphasising the importance of assessing music-generative algorithms, there is a lack of evaluation tools directly focused on detecting data replication based on raw audio. Considering this research gap, the present investigation is motivated by two main questions:

- Are audio-based music similarity metrics suitable to assess data replication in AI-generated music?
- Can we propose an open model-agnostic evaluation method and tool found on diverse audio-based music similarity metrics?

Thus, this work proposes assessing the effectiveness of five music similarity metrics<sup>2</sup> (four state-of-the-art widely-used and a novel one) in estimating exact data replication in music. We review the implications of potential data replication in AI-generated music (Section 2) and present our experimental setup, including the selected music similarity metrics and specific methodology to control and estimate exact data replication (Section 3). We analyse metrics' behaviour in different music materials (Section 4.1), aiming to assess later their data replication detection sensitivity (Section 4.2). The proposed methodology is implemented in tool MiRA (Music Replication Assessment), which computes music similarity between reference and target samples to obtain global and per-pair distances (Section 5). Finally, we discuss our research's insights, limitations and future perspectives (Section 6).

By introducing MiRA tool, we advance towards the assessment of data replication in AI-generated music using similarity metrics, contributing to open evaluation methods for accessibility for researchers, developers and users. We strive to raise awareness, detect and prevent misappropriation of training sets, and hope to motivate research on these issues.

<sup>2</sup> Hereafter, music similarity metrics refer to audio-based metrics.

## 2. BACKGROUND AND RELATED WORK

### 2.1 Implications of potential data replication in AI-generated music

Music-generative AI is advancing rapidly with novel high-quality models driven by a strong push from the industry, which is encouraging a suitable environment for real-world deployment. Yet, music generation algorithms bring significant concerns regarding their ethical, social, legal and economic implications. A key challenge is the potential data replication in AI-generated music—inquiring whether a generative model extracts and copies fragments from the training data and whether AI-generated music can be considered novel and original [3, 4]. This issue is further complicated by the implications derived concerning data misuse and IP violations such as copyright infringement. Moreover, diffusion models, one of the most popular architectures for generative AI, present high risks of data replication as they have shown a tendency to memorise their training data [5–7]. In the image generation domain, Somapalli et al. [9] demonstrate instances where generated images with diffusion models contain object-level copies of their training data. Based on image retrieval frameworks, they compare generated images with training samples and detect when content has been replicated. Similarly, Carlini et al. [5] demonstrate that diffusion models memorize and reproduce images from their training data.

Memorising training data and potential IP violations is highly under-discussed in music generative models literature, despite being one of generative AI’s main negative ethical implications in the music domain [10]. However, the recently proposed music generative model *MusicalLM* [11] has been refrained from releasing due to the ethical risks and potential work replication. In addition, *MusicalLDM* [12] acknowledges potential issues linked to data replication and plagiarism and, to address them, proposes two beat-synchronous mix-up strategies for data augmentation. The exemplified initiatives underscore the relevance of considering and addressing the ethical implications of these algorithms.

### 2.2 Evaluation methodologies in music generation

Xiong et al. [13] present a survey on music generation evaluation methodologies divided into objective, subjective and combined approaches. They highlight a current claim in finding a standardised proper method that aligns with all stakeholders, from developers to musicians and music listeners. However, even if multiple evaluation methodologies exist for music generation models, the literature highlights a lack of evaluation methodologies focused on assessing data replication and the originality of AI-generated music [4, 14]. In the symbolic domain, Yin et al. [4] introduce the *originality score* to measure the extent to which an algorithm might be copying from the training set. Nonetheless, there is a growing interest in models outputting directly audio music instead of symbolic representations. Thus, a research gap exists in detecting data replication in AI-generated music based on raw audio.

A recent work by Barnett et al. [15] proposes a framework based on two music audio embeddings to assess the similarity between the training data and AI-generated samples for understanding training data attribution. Their approach, based on VampNet [16], computes cosine distance on embeddings obtained from CLMR (Contrastive Learning of Musical Representations) [17] and CLAP (Contrastive Language-Audio Pretraining) [18].

Our perspective is that combining metrics based on audio embeddings, acoustic qualities, and features capturing music characteristics, such as chord progression or tonal similarity, provides a comprehensive assessment of potential data replication in AI-generated music. In this study, we aim to validate the effectiveness of five music similarity metrics and build an open tool to assess exact data replication in AI-generated music using these metrics.

## 3. FORCED-REPLICATION EXPERIMENT

### 3.1 Audio Music Similarity Metrics

For this study, we consider five music similarity metrics: four state-of-the-art approaches and a novel one, covering a diversity of characteristics. We here describe the metrics (summarised in Table 1) and methods used to implement them.<sup>3</sup>

**Cover Song Identification (CoverID)** [19–21]: Cover song identification is a task aiming to detect whether two music recordings are based on the same composition, accounting for variations in tempo, structure, and instrumentation while keeping a similar melodic or harmonic line. CoverID relies on pitch-content features and local alignment. To obtain CoverID distance, we use the implementation available in Essentia.<sup>4</sup> A low CoverID value suggests substantial composition similarity between the two analysed music samples.

**Kullback-Leibler (KL) divergence:** This metric provides a non-symmetric statistical measurement between reference and target probability distributions relative to their entropy. KL divergence has been employed in the literature to estimate similarity in music (e.g. [22, 23]), and more recently, to assess automatic music generation prompt adherence (e.g. [24]). We aim to explore its capabilities to estimate data replication in music samples. To obtain probability distributions, we use the PaSST audio classifier proposed in Koutini et al. [25], trained on Audioset. This methodology aligns with common practice in the literature, such as in *AudioGen* [26] and *MusicGen* [27] to obtain the probabilities of the labels in their audio and music samples. To avoid the non-symmetry of KL divergence, we compute reference to target and target to reference KL divergence and, subsequently, average both results to obtain symmetric KL divergence. Low KL divergence indicates a closer similarity between distributions.

<sup>3</sup> Two of the metrics rely on Essentia implementation. Essentia is an open-source library and tools for audio and music analysis, description and synthesis, developed in the Music Technology Group at Universitat Pompeu Fabra: <https://essentia.upf.edu>.

<sup>4</sup> [https://essentia.upf.edu/reference/std\\_CoverSongSimilarity.html](https://essentia.upf.edu/reference/std_CoverSongSimilarity.html)

**Table 1:** Summary of the considered music similarity metrics, indicating whether values correspond to higher or lower similarity ( $\downarrow/\uparrow$ ).

Metric	Description
<b>CoverID</b> ( $\downarrow$ )	Musical composition similarity based on music-specific characteristics.
<b>KL divergence</b> ( $\downarrow$ )	Differences in distributions from an audio classifier.
<b>CLAP</b> ( $\uparrow$ )	Distance between embeddings from a music pre-trained model.
<b>DEfNet</b> ( $\uparrow$ )	Novel metric based on distance between embeddings from a contrastive learning model for music similarity.
<b>FAD</b> ( $\downarrow$ )	Distance between embeddings based on CLAP music model.

**Contrastive Language-Audio Pretraining (CLAP) score** [18]: CLAP embeddings<sup>5</sup> allow to obtain latent representations of audio or text by conditioning information. For instance, *MusicLDM* [12] uses this metric to assess the novelty in text-to-music generations. To compute the CLAP score between two music samples, we extract the audio embeddings from the pre-trained music model<sup>6</sup> for each one and compute the cosine distance between them. A high CLAP score indicates a high similarity between the two music samples.

**Discogs-EffNet (DEfNet) score:** In addition to state-of-the-art distances between audio embeddings, we incorporate a novel approach based on Essentia models [28]. Essentia’s Discogs-EffNet model<sup>7</sup> provides music audio embeddings trained on Discogs metadata with contrastive learning purposes for music similarity. We consider DEfNet score to observe the effectiveness of embeddings of a model trained for a music-related task on estimating data replication. Embeddings are extracted based on track self-supervised annotations<sup>8</sup> and compute the cosine distance between reference and target samples. A high DEfNet score reveals high track similarity.

**Fréchet Audio Distance (FAD)** [29, 30]: FAD is an adaptation of Fréchet Inception Distance (FID) for music, comparing embedding distributions of a reference and a target set, based on the ViGGish model [31]. Nonetheless, a recent study by Gui et al. [30] questions whether ViGGish is the optimal model for FAD computation for music generation evaluation. They propose a tool kit<sup>9</sup> with multiple models to obtain more accurate embeddings to assess AI-generated music when calculating FAD. Consequently, we implement the adapted version of FAD using the CLAP audio music pre-trained model. A low FAD score indicates a high resemblance between the compared music samples.

<sup>5</sup> <https://github.com/LAION-AI/CLAP>

<sup>6</sup> Checkpoints: music\_audioset\_epoch\_15\_esc\_90.14.pt.

<sup>7</sup> <https://essentia.upf.edu/models.html#discogs-effnet>

<sup>8</sup> Embeddings extracted with weights `discogs_track_embeddings-effnet-bs64-1.pb`.

<sup>9</sup> <https://github.com/microsoft/fadtk>


### 3.2 Experimental Approach

To validate the effectiveness of the selected music similarity metrics in detecting exact data replication, we carried out a controlled forced-replication experiment with synthetic data, i.e. replicating music excerpts into another song under controlled conditions. Synthetic data guaranteed that the analysed music samples contained copied instances, limiting our scope to exact data replication.

For this experiment, we use an in-house dataset of 30-second audio previews from the Spotify API<sup>10</sup>, composed of over 18,000 samples and 24 music genre classes. We focus on six music genre classes defined by Spotify API internal class labels: *heavy metal*, *afrobeats*, *techno*, *dub*, *cumbia* and *bolero*. These genres were chosen for their diverse musical compositions and elements, allowing us to examine the metrics across multiple scenarios. This selection was supported using ChatGPT, which affirmed that these genres have distinct musical characteristics.

We divide data into three groups: (1) **reference set**: acting as training data, (2) **target set**: composed of synthetic data, representing AI-generated music, and (3) **mixture set**: containing different songs from the reference set but from the same music genre to build synthetic data. Synthetic data with replication contains a controlled percentage of copy from a song in our reference set: 5% (1.5s), 10% (3s), 15% (4.5s), 25% (7.5s) and 50% (15s). A synthetic sample is created by introducing the copied proportion at a random point of a music sample in the mixture set. We create 10 samples with a proportion of replication per song in the reference set. Figure 1 illustrates the procedure to build synthetic data with 5% of replication. For each music genre, the reference and mixture sets are composed of 400 songs each. Thus, the target set comprises 4,000 (400 x 10) songs per percentage of replication for each genre. Music samples are 30 seconds long as currently it is the common length in full song composition music generative models.

We assess each metric for all the songs within the reference set against themselves to establish a baseline (400 x 400 = 160,000 per-pair evaluations). Then, we compute them for each reference song and its copied instances to only consider cases with exact data replication (4,000 per-pair evaluations). Our experiment considers 120,000 samples of synthetic data (approximately 167h of music with a proportion of data replication).



**Figure 1:** Synthetic data procedure with 5% of replication.

<sup>10</sup> <https://developer.spotify.com/documentation/web-api>

## 4. RESULTS

### 4.1 Analysing metric behaviour

Figures 2, 3, 4, 5 and 6 depict the average  $\mu$  and standard deviations  $\sigma$  of the different metrics per degree of replication and music genre. We observe a steady and similar behaviour by three metrics (CoverID, CLAP and DEfNet) through all studied music genres, showing higher similarity values for cases with higher replication levels (i.e., 50%). Standard deviation decreases with increasing replication level, which suggests less disparity within the analysed pairs. These three metrics show the sensitivity<sup>11</sup> to estimate data replication.

Instead, KL divergence presents a different behaviour with very similar values of  $\mu$  and  $\sigma$  for all degrees of replication. Some sensitivity is observed in all music genres, except for *dub*, where the baseline mean  $\mu_b$  is smaller than in replication cases  $\mu_r$ , despite the standard deviation being higher ( $\mu_b=0.757$ ,  $\sigma_b=0.511$ ;  $\mu_r=0.862$ ,  $\sigma_r=0.462$ ). Thus, KL divergence demonstrates the capability of detecting replication but is ineffective in distinguishing between degrees of replication.

Contrasting with the other metrics, FAD based on CLAP music embeddings completely differs from them. On the one side, its behaviour is inconsistent as it exhibits fluctuating trends for the different examined cases. On the other side, it fails to detect data replication. A higher similarity value (low FAD) is always obtained for the baseline. Instead, for the different degrees of replication, higher FAD is achieved. Consequently, FAD based on CLAP music embeddings does not appear to be a suitable metric to assess exact data replication in music samples.

By analysing the metrics' behaviour, we could directly conclude that CoverID, KL divergence, CLAP and DEfNet are suitable for our posed research aim. However, further exploration is required before determining their ability to detect replication and degree of replication. We delve into this analysis in the next subsection.

### 4.2 Assessing data replication detection sensitivity

In this section, we complement the previous analysis with an assessment of statistical differences. Because our data is not normally distributed and variance is heterogeneous, the Kruskal-Wallis test [32] is the most adequate statistical analysis to examine our results, as is non-parametric, does not rely on normality and handles unequal sample sizes. We perform the Kruskal-Wallis test on CoverID, KL divergence, CLAP and DEfNet. Significant statistical differences ( $p < 0.05$ ) are observed across all music genres and degrees of replication, consistent with our earlier findings.

Nonetheless, the insight of this analysis relies on the pairwise comparisons between the baseline and different degrees of replication. CoverID pairwise comparison reveals a statistically significant difference between the baseline and the 5% replication degree for *afrobeat*, *cumbia* and *techno*. For the three other music genres, this happens




Figure 2: CoverID (↓)




Figure 3: KL divergence (↓)




Figure 4: CLAP (↑)

<sup>11</sup> Sensitivity is understood as the capability to differentiate between degrees of replication.




Figure 5: DfNet (↑)




Figure 6: FAD (↓)

for a 10% replication degree. Then, statistical significance also appears in pairwise comparisons of different degrees of replication. We can derive that CoverID is sensible for 10% of replication, and in some cases at 5%. When considering KL divergence, pairwise comparison depicts a statistically significant difference between the baseline and the 5% replication degree. Between degrees of replication, no statistical significance is revealed for any pairwise comparison, except for *heavy metal* between 5% and the other replication degrees. Regarding the CLAP and DfNet, a significant difference already appears when comparing the baseline against the samples with 5% replication, indicating that these metrics are sensitive to 1.5 seconds of replication. In all cases, a notable difference emerges among the levels of replication, enhancing the sensitivity of these metrics' detection capabilities. They demonstrate sensitivity to varying replication degrees.

Withal, this statistical analysis sustains the validity of these four metrics to assess exact data replication in the training set and determines their degree of sensitivity.

## 5. MUSIC REPLICATION ASSESSMENT TOOL

Derived from the presented experiment, we implement the proposed methodology into an evaluation tool. We introduce the Music Replication Assessment (MiRA) tool: an open evaluation method based on four diverse raw audio music similarity metrics.

MiRA computes music similarity between reference and target samples to obtain global and per-pair distances, based on CoverID, KL divergence, CLAP and DfNet. It can estimate data replication with a proportion higher than 10% (3 seconds), but in most of the examined scenarios, it is sensible to 5% of replication. Per-pair distances are highly beneficial for detecting close pairs, outliers and suspicious cases with potential data replication. Considering that replication detection requirements may vary depending on the evaluation, users are left to set their replication threshold. In addition, MiRA is model-independent as no information about the model architecture or its characteristics is necessary. The evaluation is conducted directly with the training (**reference**) and generated samples (**target**) of the analysed generative model.

However, designating a baseline value is encouraged to accurately interpret the music similarity between the reference and target samples. We propose a third comparison group of samples (**control**) based on songs related to the reference songs but unseen by the model (e.g. shared music genre). Again, this is a decision for the users conditioned to their evaluation scope. Note that using a control group allows us to understand and interpret the results obtained by acting as the baseline similarity level of independent songs with a shared characteristic.

The complete structure of the implemented system is depicted in Figure 7. We release MiRA as an open-source tool, built into a PyPI package<sup>12</sup>. Together with the code, we provide examples and best practice recommendations for using this methodology. With the release of MiRA, we hope to enhance transparency in music generation models and data replication assessment.

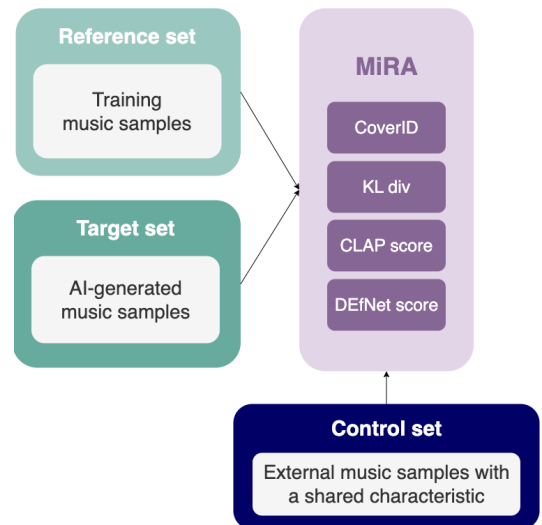


Figure 7: MiRA's structure scheme.

<sup>12</sup> <https://pypi.org/project/mira-sim/>

## 6. DISCUSSION AND CONCLUSIONS

This work focused on validating the use of music similarity metrics for assessing data replication in AI-generated music. We hypothesise that similarity metrics are effective in estimating data replication. Therefore, we framed the scope of our study to exact data replication in music samples, while conducting a controlled forced-replication experiment with synthetic data.

We examined five diverse audio-based metrics: four standard metrics (CoverID, KL divergence, CLAP and FAD) and a novel one (DefNet). Our results indicate that four of the five studied metrics can detect data replication to a certain extent. Instead, FAD based on CLAP music embeddings presented an opposite behaviour compared to the other metrics. Higher similarity is obtained for the baseline group and FAD shows unstable trends throughout the diverse music genres. Thus, we do not find it suitable for our case study. However, it must be acknowledged that the recent publication by Gui et al. [30] offered multiple classifiers to compute FAD. There is the possibility that we did not consider the appropriate classifier for our task. Thus, we should consider exploring other classifiers before determining the validity of FAD in detecting replication in music.

Regarding the other four metrics, our results show interesting insights. First, we find CoverID to be sensible to different replication degrees, establishing a robust threshold level at 10% of replication. Furthermore, in some of the studied cases, replication sensitivity is lowered to 5% of replication. This is a substantial finding to validate the suitability of metrics oriented to specific music characteristics, such as tempo, structure and composition.

Next, we observe that KL divergence can be sensitive to replication as pairwise comparison between baseline and degrees of replication is statistically significant. Nevertheless, the other pairwise results reveal that KL divergence is ineffective for differentiating between replication degrees. We consider this an unexpected turnout in our analysis.

Considering CLAP and DefNet scores, both embedding-based metrics, our experiment validates their suitability to detect data replication. Not only do they show robustness by increasing their similarity value parallel to the replication degrees (i.e. higher similarity for higher level of replication), but they also show high sensitivity for different degrees of replication. All results suggest their sensitivity might be higher than we envisioned and might be able to detect replication in smaller samples (i.e. < 1.5 seconds).

As a result of these findings, we achieve our second goal within the scope of this research: to build an open model-agnostic tool based on music similarity metrics on raw audio. In this article, we have introduced the MiRA tool, leveraging the four validated similarity metrics, which can be used to evaluate any music-generative model with audio output. MiRA does not require any information about the model architecture or its characteristics. Instead, similarity evaluation relies on comparing reference and target samples.

By introducing the MiRA tool, we are contributing to the research gap of lack of evaluation methodologies directly assessing potential data replication in AI-generated music. Our study validates the use of similarity metrics to estimate training data replication. We intend to encourage the open evaluation of music generation models by researchers, developers and users concerning data replication. In addition, our research strives for the importance of ethical, social, legal and economic consequences of generative AI in the music domain, together with the need to address their risks and issues.

### 6.1 Limitations and Future Work

Despite our contribution to advance towards data replication assessment with music similarity metrics, there are multiple opportunities to complement our investigation.

First, we limited the scope of our experimental approach to assessing the use of different music similarity metrics for exact data replication, consequently reducing the definition of plagiarism to exact replication of fragments in the training set. We followed such an approach to validate our hypothesis and ensure an attainable method to address this issue. While this reduced scope could potentially be solved using audio fingerprinting strategies [33], we believe that by employing a diverse range of metrics we can provide a more comprehensive assessment of data replication.

Framing our aim to exact data replication also introduced a limitation in considering typical perturbations that music samples experience when training the model or during the model procedure to generate a music sample. Thus, it would be a key point for future work to validate the robustness of these metrics towards typical data augmentation techniques, such as pitch shifting and reverberation. Proving them to be robust would also enhance the capabilities of MiRA for detecting potential replication in AI-generated music. At the same time, we intend to expand the abilities of MiRA for data replication by incorporating complementary metrics, if necessary.

In addition, our experimental process was limited to the high computational costs of some of the metrics. In particular, we faced significantly large amounts of time to compute FAD and KL divergence. This is a relevant concern as we want MiRA to be an open tool that can be used by any researcher or user. Thus, considering the computational capacity required to compute the integrated metrics within is a relevant issue in our research.

Another limitation is the type of data that we use. We base our experiment on synthetic data despite our goal being oriented to AI-generated music. We must use synthetic data with a controlled percentage of replication to guarantee and assess the capabilities of detection and sensitivity of music similarity metrics. However, we would like to test the validity of the introduced tool when used in a generation context. To do so, we require not only a generative model but its details on training data and generation samples. We plan to expand our research in with AI-generated content in upcoming studies.

## 7. ACKNOWLEDGMENTS

The authors would like to thank Gaëtan Hadjeres and William Thong from Sony AI, as well as Dmitry Bogdanov and Pablo Alonso-Jiménez from the Music Technology Group at Universitat Pompeu Fabra, for their insightful discussions and valuable feedback throughout the development of this research.

## 8. ETHICS STATEMENT

The late rapid popularity growth of generative AI in the music domain brings significant ethical implications. The main challenges are linked to the role of AI within music creative processes, such as composition, potential misappropriation of data in AI-generated music, inquiries on the novelty of generations, derived authorship attribution, effects on intellectual property rights and sustainability of current business models. In addition, there are notable concerns about the cultural bias in these systems and their environmental impact.

Our research focused on the issue of assessing potential data replication in AI-generated music. We observed a lack of evaluation methodologies to examine replication in raw audio. We contributed to this issue by proposing a methodology based on audio-based music similarity metrics. We demonstrated its effectiveness and provided an open tool to evaluate AI-generated music. Our introduced approach is contributing to the transparency of music generation algorithms.

Despite the positive contribution of our investigation, we must be critical of some methodological aspects of our work. Our principal ethical concern falls under the type of data used to conduct our forced-replication experiment. In particular, we employ an internal dataset created with Spotify previews (30-second samples of music). Even if these practices are common in the ISMIR community, we see the need for guidelines for the legal assessment of MIR data included in datasets, incorporating country dependencies, origin and intended use, personal data involved (from artists and listeners) and potential future consequences<sup>13</sup>.

## 9. REFERENCES

- [1] F. Carnovalini and A. Rodà, “Computational creativity and music generation systems: An introduction to the state of the art,” *Frontiers in Artificial Intelligence*, vol. 3, 2020.
- [2] E. Gómez, M. Blaauw, J. Bonada, P. Chandna, and H. Cuesta, “Deep learning for singing processing: Achievements, challenges and impact on singers and listeners,” *ArXiv*, 2018. [Online]. Available: <https://arxiv.org/abs/1807.03046v1>
- [3] B. L. T. Sturm, M. Iglesias, O. Ben-Tal, M. Miron, and E. Gómez, “Artificial intelligence and music: Open questions of copyright law and engineering praxis,” *Arts*, vol. 8, p. 115, 2019.
- [4] Z. Yin, F. Reuben, S. Stepney, and T. Collins, “Measuring when a music generation algorithm copies too much: The originality report, cardinality score, and symbolic fingerprinting by geometric hashing,” *SN Computer Science*, vol. 3, 2022.
- [5] N. Carlini, J. Hayes, M. Nasr, M. Jagielski, V. Sehwag, F. Tramer, B. Balle, D. Ippolito, and E. Wallace, “Extracting training data from diffusion models,” in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 5253–5270.
- [6] N. Carlini, D. Ippolito, M. Jagielski, K. Lee, F. Tramer, and C. Zhang, “Quantifying memorization across neural language models,” *ArXiv*, 2023.
- [7] D. Bralios, G. Wichern, F. G. Germain, Z. Pan, S. Khurana, C. Hori, and J. L. Roux, “Generation or replication: Auscultating audio latent diffusion models,” *ArXiv*, 2023.
- [8] H. Wang, “Authorship of artificial intelligence-generated works and possible system improvement in china,” *Beijing Law Review*, vol. 14, pp. 901–912, 2023.
- [9] G. Somepalli, V. Singla, M. Goldblum, J. Geiping, and T. Goldstein, “Diffusion art or digital forgery? investigating data replication in diffusion models,” *ArXiv*, 2022.
- [10] J. Barnett, “The ethical implications of generative audio models: A systematic literature review,” *AIES 2023 - Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 146–161, 2023.
- [11] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzetti, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi, M. Sharifi, N. Zeghidour, and C. Frank, “MusicLM: Generating music from text,” *ArXiv*, 2023.
- [12] 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,” *ArXiv*, 2023.
- [13] Z. Xiong, W. Wang, J. Yu, Y. Lin, and Z. Wang, “A comprehensive survey for evaluation methodologies of AI-generated music,” *ArXiv*, 2023.
- [14] R. Battle-Roca, E. Gómez, W. Liao, X. Serra, and Y. Mitsufuji, “Transparency in music-generative AI: A systematic literature review,” *Research Square preprint*, 2023.
- [15] J. Barnett, H. F. Garcia, and B. Pardo, “Exploring musical roots: Applying audio embeddings to empower influence attribution for a generative music model,” *arXiv*, 2024.

<sup>13</sup> We refer to a recently documented example of research vs legal clash linked to algorithmic auditing in the music domain <https://www.rollingstone.com/pro/features/spotify-teardown-book-streaming-music-790174/>

- [16] H. F. Flores Garcia, P. Seetharaman, R. Kumar, and B. Pardo, "Vampnet: Music generation via masked acoustic token modeling," in *Proceedings of the 24th International Society for Music Information Retrieval Conference, ISMIR 2023, Milan, Italy, 2023*.
- [17] J. Spijkervet and J. A. Burgoyne, "Contrastive learning of musical representations," *Proceedings of the 22nd International Society for Music Information Retrieval Conference, ISMIR 2021, Online, 2021*.
- [18] Y. Wu, K. Chen, T. Zhang, Y. Hui, T. Berg-Kirkpatrick, and S. Dubnov, "Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation," *ArXiv, 2023*.
- [19] J. Serrà, X. Serra, and R. Andrzejak, "Cross recurrence quantification for cover song identification," *New Journal of Physics*, vol. 11, 2009.
- [20] J. Serrà, E. Gómez, P. Herrera, and X. Serra, "Chroma binary similarity and local alignment applied to cover song identification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 6, pp. 1138–1151, 2008.
- [21] J. Serrà, E. Gómez, and P. Herrera, "Transposing chroma representations to a common key," *IEEE CS Conference on The Use of Symbols to Represent Music and Multimedia Objects, 2008*.
- [22] M. Hoffman, D. Blei, and P. Cook, "Content-based musical similarity computation using the hierarchical dirichlet process." in *Proceedings of the 9th International Society for Music Information Retrieval Conference, ISMIR 2008, Philadelphia, USA, 2008*.
- [23] D. Schnitzer, A. Flexer, G. Widmer, and M. Gasser, "Islands of gaussians: The self organizing map and gaussian music similarity features," in *Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR 2010, Utrecht, Netherlands, 2010*.
- [24] Z. Evans, C. Carr, J. Taylor, S. H. Hawley, and J. Pons, "Fast timing-conditioned latent audio diffusion," *ArXiv, 2024*.
- [25] K. Koutini, J. Schlüter, H. Eghbal-zadeh, and G. Widmer, "Efficient training of audio transformers with patchout," in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea. ISCA, 2022*, pp. 2753–2757.
- [26] F. Kreuk, G. Synnaeve, A. Polyak, U. Singer, A. Défossez, J. Copet, D. Parikh, Y. Taigman, and Y. Adi, "Audiogen: Textually guided audio generation," 2023.
- [27] J. Copet, F. Kreuk, I. Gat, T. Remez, D. Kant, G. Synnaeve, Y. Adi, and A. Défossez, "Simple and controllable music generation," *ArXiv, 2023*.
- [28] P. Alonso-Jiménez, D. Bogdanov, J. Pons, and X. Serra, "Tensorflow audio models in Essentia," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2020*.
- [29] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi, "Fréchet Audio Distance: A Reference-Free Metric for Evaluating Music Enhancement Algorithms," in *Proc. Interspeech 2019, 2019*, pp. 2350–2354.
- [30] A. Gui, H. Gamper, S. Braun, and D. Emmanouilidou, "Adapting frechet audio distance for generative music evaluation," *ArXiv, 2023*.
- [31] S. Hershey, S. Chaudhuri, D. P. W. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, M. Slaney, R. J. Weiss, and K. Wilson, "Cnn architectures for large-scale audio classification," *ArXiv, 2017*.
- [32] W. H. Kruskal and W. A. Wallis, "Use of ranks in one-criterion variance analysis," *Journal of the American Statistical Association*, vol. 47, no. 260, pp. 583–621, 1952.
- [33] P. Cano and E. Batlle, "A review of audio fingerprinting," *Journal of VLSI Signal Processing*, vol. 41, pp. 271–284, 11 2005.