

A DATASET AND BASELINES FOR MEASURING AND PREDICTING THE MUSIC PIECE MEMORABILITY

Li-Yang Tseng Tzu-Ling Lin Hong-Han Shuai
Jen-Wei Huang Wen-Whei Chang

National Yang Ming Chiao Tung University, Hsinchu, Taiwan

{liyangtseng.ee10, tzulinglin.11, hhshuai, admsd.ee10, wwchang}@nycu.edu.tw

ABSTRACT

Nowadays, humans are constantly exposed to music, whether through voluntary streaming services or incidental encounters during commercial breaks. Despite the abundance of music, certain pieces remain more memorable and often gain greater popularity. Inspired by this phenomenon, we focus on measuring and predicting music memorability. To achieve this, we collect a new music piece dataset with reliable memorability labels using a novel interactive experimental procedure. We then train baselines to predict and analyze music memorability, leveraging both interpretable features and audio mel-spectrograms as inputs. To the best of our knowledge, we are the first to explore music memorability using data-driven deep learning-based methods. Through a series of experiments and ablation studies, we demonstrate that while there is room for improvement, predicting music memorability with limited data is possible. Certain intrinsic elements, such as higher valence, arousal, and faster tempo, contribute to memorable music. As prediction techniques continue to evolve, real-life applications like music recommendation systems and music style transfer will undoubtedly benefit from this new area of research.

1. INTRODUCTION

Music memorability is essential and has a wide range of commercial applications. For instance, content creators and marketing teams can use unique visual aids or audio components to captivate target audiences and distinguish themselves from other information sources [1, 2]. Sound logos, such as Netflix’s iconic “ta-dum,” are designed to engage listeners and promote brand recognition. In the realm of cognition literature, numerous studies have sought to understand the factors that contribute to music memorability [3–6]. For instance, [5, 6] bridged the gap between cognitive science and MIR by examining whether implicit or explicit memory for a single tune is impacted by

the type of encoding task and variations in timbre, tempo and structure.

However, music memorability remains a relatively unexplored area, particularly from a data-driven standpoint. Research related to music memorability includes the study of involuntary musical imagery (INMI) [7, 8], also known as “earworms,” which refers to the phenomenon where fragments of music become mentally lodged on repeat. For instance, Jakubowski et al. proposed a model that can determine whether a piece of music may induce the INMI effect by using statistical analysis and a random forest model [8]. However, the mechanism of INMI differs from music memorability since the former is a passive process while the latter can be active, e.g., everyone remembers how to sing “Happy Birthday,” but the song may not qualify as an earworm. Another line of prior studies [9–12] investigating the intrinsic memorability of multimedia content have predominantly focused on computer vision, with their findings suggesting that data-driven approaches can effectively determine memorability levels. Motivated by these studies, we break new ground in exploring music memorability from a data-driven perspective by compiling a novel dataset and employing machine learning techniques.

Specifically, to expand the scope of memorability detection and recognition in music information retrieval (MIR), we establish a new research domain called music memorability regression (MMR), which aims to predict a memorability score for a given music piece. We create an experimental procedure as shown in Figure 1 to collect a new dataset, the YouTube Music Memorability (YTMM) dataset, where memorability scores are determined by the percentage of participants who can recall the music piece after a certain period. This dataset provides reliable and consistent music memorability scores across all participants, paving the way for further research in the field. We also propose several baseline approaches for predicting music memorability, including feature engineering using hand-crafted music-related features and transfer learning techniques. These baselines not only demonstrate the potential of machine learning in addressing music memorability but also serve as a foundation for future work.

Despite the promise of machine learning in tackling music memorability by predicting memorability scores, its “black box” characteristics hinder the interpretation of machine decisions in MIR tasks. A straightforward approach would be to compute correlations without relying on black-



© L.-Y. Tseng, T.-L. Lin, H.-H. Shuai, J.-W. Huang, and W.-W. Chang. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** L.-Y. Tseng, T.-L. Lin, H.-H. Shuai, J.-W. Huang, and W.-W. Chang, “A dataset and Baselines for Measuring and Predicting the Music Piece Memorability”, in *Proc. of the 24th Int. Society for Music Information Retrieval Conf.*, Milan, Italy, 2023.




Figure 1. The music memory game, which allows data annotators to label music memorability scores reliably. The experiment is divided into three stages, each with a 3-minute long break in between. Each 18-minute stage is composed of multiple 5-second music pieces and short breaks.

box prediction models to glean insights about the relationship between memorability and musical features. However, given the complexity of analyzing music memorability, using a single feature results in an extremely low correlation with memorability, leading to inconclusive findings. One alternative would be to explore all possible feature combinations when calculating correlations, but the sheer number of combinations, e.g., $2^{20} - 1$ for just 20 features, renders this approach impractical. A/B testing could be used to determine which type of music is more memorable, but it suffers from similar drawbacks, such as being time-consuming and unable to account for all variables that may impact the experiment’s outcome. To make machine learning models reveal their “black box” characteristics, researchers are increasingly adopting explainable artificial intelligence (XAI) [13] for deeper insights. Building on previous interpretability analyses in audio processing [14, 15], we utilize Shapley Additive Explanations (SHAP) [16], a game-theoretic approach that clarifies the output of machine learning models, to identify the key components of memorable music.

Our main contributions are as follows: first, we present the new YTMM dataset with objective annotations of memorability scores, which will be publicly available for future research; second, we propose several deep learning baseline models for MMR; and finally, we explore the potential characteristics of memorable music pieces while providing interpretability for these deep learning-based methods.

2. RELATED WORK

In addition to the cognition literature on music memorability [5, 6], there are several related yet distinct terms, such as Involuntary Musical Imagery (INMI) or “earworms”—fragments of music that involuntarily come to mind [7]. Studies have examined earworms through interviews, environmental and psychological conditions leading to INMI, and the impact of melodic features and song popularity on spontaneous musical imagination [8]. Crucial differences between INMI and music memorability include: 1) INMI involves uncontrollable mental repetition, while memorability requires conscious recall; and 2) the stimuli in [8] are highly familiar to participants, whereas our study selects audios unfamiliar to most annotators to

mitigate the influence of individual listening histories on memorability. Another related concept is hook catchiness [17–20], which refers to the most easily recalled fragment of a musical piece. However, our focus lies in predicting the memorability of different music pieces rather than assessing the impact of various segments within the same tune on catchiness prediction and recognition. Furthermore, we ensure our stimuli consist solely of pure instrumental music clips to prevent any textual information from lyrics influencing music memorability.

Moreover, while deep learning has achieved significant success in supervised MIR tasks, it often demands large-scale annotated data. However, collecting useful annotations for MIR tasks can be costly, as it typically requires expertise and domain knowledge [21]. To tackle this challenge, various data augmentation and training strategies have been proposed [21–24]. For instance, McFee et al. [21] apply transformations such as pitch shifting, time-stretching, and adding background noise to the original waveform. Cubuk et al. [22] mask both time and frequency content to expand the input space in automatic speech recognition (ASR) and MIR tasks. To enhance learning robustness with limited data, Wu et al. [23] extract general music representations using a multi-task pre-trained encoder, inspired by speech processing research [25, 26]. Similarly, Castellon et al. [24] employ transfer learning from existing music generation architectures. However, not all the aforementioned methods are simultaneously open-source, computationally inexpensive, and interpretable. Therefore, in this paper, we focus on applying signal processing approaches like masking, with further details provided in Section 4.


3. DATASET CONSTRUCTION FOR MUSIC MEMORABILITY

In this section, we discuss the details of our dataset collection process and how music memorability is measured.

3.1 Audio Collection

To construct a dataset with objective music memorability scores, we first ensure that the audio samples are unbiased. We randomly select music by querying music-related videos using the YouTube API with random query keys, avoiding any specific music genre preference. Addi-

Task Type	# of Audios	# of Repetition (min)	# of Repetition (max)	# of Repetition (avg)	# of Repetition (std)
Filler	65	-	-	-	-
Vigilance	21	5	10	6.5	1.08
Short-Term Target	88	10	49	25.23	10.81
Medium-Term Target	41	61	131	110.33	16.32
Long-Term Target	20	155	276	222.05	36.64

Table 1. Details of different audio tasks in the music memory game.

Figure 2. Distributions of the audio published location and the distributions of the audio views in the final dataset.

tionally, we manually filter the music to confirm that the queried videos contain pure music content, excluding instrument tutorials or gadget unboxings. Next, we conduct a pilot study to verify that the selected audios are unfamiliar to most of the annotators in our target user group. Considering the annotators’ nationalities might not be as varied as the music collection’s, and language can be a memorable yet non-music-related element, we only use the intro part of each song. This approach helps eliminate other potential variables affecting music memorability. Also, the volume across all audio clips is normalized to minimize any memorable attributes unrelated to the music itself. Loudness normalization ensures the music is remembered based on its inherent qualities rather than its loudness.

We use only a segment of each audio for two reasons: i) to better eliminate confounding factors, such as vocal timbre, and ii) to shorten the period of annotations and prevent fatigue. We achieve this by truncating audios into structurally meaningful segments and applying proper time-stretching to alter the duration of an audio signal to a fixed length without distorting the audio. The segmentation process is supervised by an expert with a professional music education background. Note that time-stretching not only reduces modeling complexity but also prevents annotators from memorizing the audio based on its duration.

Ultimately, we collect 235 structurally meaningful 5-second audios with labeled music memorability scores. Our goal is to determine which types of music pieces are more likely to be memorized, rather than focusing on entire music clips, which are more complex and involve additional factors. This research can facilitate various applications, such as Netflix’s iconic “ta-dum” sound. The collected data can be found in the supplementary materials. Figure 2 illustrates the distribution of the collected audios concerning their published geographical locations and total views on YouTube, with view counts ranging from 10K to 100M.

3.2 The Music Memory Game

To annotate the memorability of the collected musical data, we follow the setting of image memory game [9] to design a novel music listening experiment. During the experiment, the recruited data annotators are asked to listen to 235 music pieces in total and answer whether the audio is repeated in the experiment or not. From a cognitive view, we define music memorability as long-term musical salience and the extent to which a musical piece continues to be remembered over time. In the music memory game, music memorability is measured as the tendency to correctly recognize a music piece when encountering it again in the experiment among all users. Specifically, let $x_j^{(i)}$ denote whether the i -th music piece can be recalled by the j -th data annotator, *i.e.*, 1 if the annotator recognized the i -th music piece. The memorability score of music i , denoted by $m^{(i)}$, is then calculated by:

$$m^{(i)} = \frac{1}{n_i} \sum_{j=1}^{n_i} x_j^{(i)}, x_j^{(i)} \in \{0, 1\} \quad (1)$$

where n_i is the total number of data annotators for the i -th music.

To make the ground truth unbeknownst to all participants, music excerpts are split into three task categories: “vigilance”, “target”, and “filler”. Targets and vigilance targets are both repeated in the experiment, while the former are collected to be the true labels and the latter is used to make sure participants are attentive when labeling data. Moreover, fillers are used to stuff the spacing between the first and second repetition of a target and therefore is only presented once. The overview of the music memory game experimenting procedure is shown in Figure 1. The target-vigilance-filler split details can be found in Table 1. Rigorous criteria are enforced to monitor the performances of data annotators and preserve the quality of memorability labels. Specifically, annotations from users who detect vigilance repetition with an accuracy lower than 60% are automatically discarded. Furthermore, to prevent gathering biased memorability, all annotators only engage in labeling once. We recruited a total of 218 users from campus, with 163 clearing the vigilance accuracy level, 17% of passed annotators having professional music education backgrounds, and over 98% being between the ages of 20 and 29.

Differing from previous works on image memorability, our experiment is composed of three similar stages with breaks inserted in between. The reasons for using stages and breaks are two-fold. First, audios are sequential, there-




Figure 3. Memorability scores at various stages. The color symbolizes the rank of short-term memorability, while the lines represent stage relationships. The plot also shows Spearman’s rank correlations ρ between memorabilities measured at each stage.

fore it is more exhausting to label the memorability score to audios as compared to static images. Second, it usually takes some time for the earworm phenomenon to happen when listening to music. Hence, we assume memorability should be invariant even after encountering breaks that probably would reset the memory. The results of relations between repeat interval and memorability score are shown in Figure 3, where the lines exhibiting memorability scores across short-term, medium-term, and long-term repeats. The results manifest that the memorability score is indeed independent of the sequential context. Therefore, it is easier to memorize truly memorable pieces of music even after long breaks. The fact that Spearman’s rank correlation [27] between short-term, medium-term, and long-term are all greater than 0.64 also proves that the rank of memorability score is preserved across variant repeat intervals.

3.3 Labels and Consistency Analysis

To assure that collected labels are universal across all data annotators, we evaluate the human consistency according to previous work [9] by randomly splitting all participants into 2 groups and examining how well the memorability scores measured in the first groups matched the ones measured in the second group by averaging Spearman’s rank correlation [27] between randomly separated two halves of the participants 25 times. The average Spearman’s rank correlation coefficient ρ is 0.83, indicating the consistency of the collected data.

Figure 4 shows the scattering plot of music memorability and repeat interval. The graph demonstrates that mu-




Figure 4. Relations between memorability score and target repeat interval in log scale. The hue represents the level of fatigue.

sic possesses a linear relation between memorability score and log-scaled repeat interval. Please note that the fatigue level is another factor in the plot that also contributes to the memorability score of audio. The fatigue level, defined as the amount of audios listened without a 3-minute break, is a direct result caused by staging experiment and participating in taking a break in the middle since listening to more music at one time without resting reduces participants’ ability to identify repeated music pieces. The setting of inserting audio to random positions in the experiment procedure adds more context diversity to the process of memorizing music, thus making the labeled memorability scores more robust.

4. MUSIC MEMORABILITY PREDICTION

4.1 Learning with Handcrafted Features

Although feature extractions for deep learning models can be data-driven without being handcrafted, leading to a better result given sufficient training data, handcrafted features provide interpretable information for more insights. Therefore, we propose handcrafted features that can more accurately depict the low-level acoustic features or high-level semantic features of musical clips as shown in Table 2. For the low-level acoustic features that can be directly derived from the audio signal of music segments, we leverage the harmony, rhythm and timbre since they are most easily recognizable fragments of a piece of music [17] and describe the fundamental elements of a tune. Moreover, zero crossings and zero crossing rate are also extracted since they give the impressions into the frequency content of a signal. On the other hand, high-level semantic features are more abstract descriptions. Since the previous works in Psychology [28,29] mention the link between music emotion and memory, we introduce valence and arousal, which represent the mood of music pieces as features.

Another high-level feature is genre, which describes

Level	Category	Feature Implementation
Low-level	Harmony	mean, std of 12 pitch class
	Rhythm	beat per minute (bpm)
	Timbre	mean, std of 4-tracks (Vocals, Bass, Drums, Others)
	Zero Crossing	# of zero crossings & avg, median of zero crossings rate
High-level	Mood	valence, arousal
	Genre	Music, Musical Instrument

Table 2. Explainable handcraft features.

how likely a clip is belong to a certain type of music. Specifically, due to the unstable performance of existing algorithms for detecting sequences of chord labels, we employ chromagram (chroma) [30] as a representation of harmony patterns. To extract timbre information, the Mel-Frequency Cepstral Coefficient (MFCCs) is widely utilized. Although MFCCs is representative for timbre, its components are difficult to grasp intuitively. As a result, we treat MFCCs as a raw feature and find an alternative solution by first separating source audios into four components using source separation software Spleeter [31], and calculating their respective amplitudes to represent the characteristics of different instruments and frequency ranges. For the rhythmic pattern, although Tempogram [32] captures the underlying rhythmic pattern of raw audios, it is unable to provide precise insights to concretely measure the audio’s groove. Therefore, we instead utilize beat per minute (bpm) to represent general rhythm characteristics. We also use static valence and arousal values to describe perceived music moods, which are predicted by using Support Vector Regression (SVR) with a linear kernel trained on the PMemo dataset [33]. For genre features, we use the predicted music tagging and instruments from the downstream task of PANN [34]. Finally, SVR and Multilayer Perceptron (MLP) are employed as predictors to link audio features to memorability scores.

4.2 End-to-End Deep Learning

Deep learning [35] is featured by its ability to directly learn meaningful information from raw data, as opposed to using hand-crafted features. As a result, we also test end-to-end models to find if their feature-learning process improves performance. Our model uses spectrograms in Mel-scale as inputs, similar to previous end-to-end MIR tasks. Moreover, transfer learning [36], which applies previously learned knowledge to new data, has been found to significantly increase learning performance by skipping costly data-labeling procedures. Here, we use the self-supervised pre-trained Audio Spectrogram Transformer (SSAST) [37] since SSAST has been proved to achieve state-of-the-art results on numerous audio tasks, including audio event classification, keyword spotting, mood recognition, and speaker identification, after being trained on a vast amount of unlabeled data.

Method	Corr.	MSE	MSE STD
chroma + MLP	0.1740	0.0326	-
MFCCs + MLP	0.1179	0.0353	-
convnet features [38] + MLP	0.1889	0.0314	-
EHC features + SVR	0.2988	0.0339	0.0128
EHC features + SVR + PS	0.2084	0.0391	0.0129
EHC features + MLP	0.2656	0.0263	0.0058
EHC features + MLP + PS	0.2388	0.0275	0.0059
mel-spectrograms + SSAST	0.0124	0.0298	0.0061
mel-spectrograms + SSAST + PS	0.2658	0.0265	0.0074

Table 3. Spearman’s rank correlation and MSE loss between predicted and ground truth music memorability score using different models. Note that EHC features stand for explainable handcrafted features, PS stands for pitch shift data augmentation, and Corr. represents Spearman’s rank correlation.

5. EXPERIMENT RESULTS

Evaluation Metrics. Spearman’s rank correlation and mean squared error (MSE) loss are used as the metrics to evaluate the performance of music memorability prediction. The former indicates the ability to rank the relative memorability of different audios, while the latter indicates the absolute error of the predicted results.

Different Baselines. Here, we leverage Chroma and MFCCs along with their respective derivatives as two hand-crafted feature representations and fit the ground truth by Multilayer Perceptron (MLP) as two simple baselines. Moreover, we also use the convnet model as a baseline since it is the most referenced and available work in general music representation. The convnet model [38] utilizes CNNs for music tagging in the pre-training stage, and the extracted features serve as the representation for downstream tasks. Finally, Self-Supervised Audio Spectrogram Transformer (SSAST) [37] is also used as the baseline, which is a Transformer-based model with more parameters as compared to CNNs. SSAST pretrains the model with joint discriminative and generative masked spectrogram patch modeling.

Implementation Details. All the feature classifiers are pretrained without finetuning on the self-collected dataset. To handle the instability stemming from the limited labeled data, we normalize labels by subtracting the mean value, *i.e.*, predicting a relative value instead of an absolute value. For MLP and SSAST models, the learning rates are respectively set to $5e-5$ and $5e-6$ with the Adam optimizer [39]. We also conduct additional feature selection on the handcrafted features to improve the convergence of the MLP/SVR model (only select 25 features) due the small number of data samples. In addition, techniques including frequency masking, band stop filtering, and reverberation [26] are used to augment data, together with the pitch shifting augmentation. The results are reported by the average of the 10-fold outputs.

5.1 Prediction Results.

Table 3 compares the results of different prediction models, where SVR and MLP take explainable handcrafted

Model	Top-k feature selection	Corr.	MSE
MLP	k = 40 (no feature selection)	0.2160	0.0272
MLP	k = 35	0.2324	0.0270
MLP	k = 25	0.2656	0.0263
MLP	k = 20	0.2018	0.0271
SVR	k = 40 (no feature selection)	0.2168	0.0324
SVR	k = 35	0.2291	0.0340
SVR	k = 25	0.2988	0.0339
SVR	k = 20	0.2630	0.0354

Table 4. Spearman’s rank correlation and MSE loss for MLP/SVR models with different top-k feature selection.

features as inputs, and SSAST takes mel-spectrograms as inputs. The results indicate that chroma and MFCCs produce the worst results due to the ineffective feature extraction. For convnet features, the performance is better than chroma and MFCCs due to the pretraining. However, the amount of training data is too small to finetune the model on the music memorability regression task. SSAST outperforms other baselines since it incorporates the prior knowledge of spectrograms pre-trained by using advanced methods. Finally, Explainable Handcrafted Features (EHC) method produces the best correlation results by combining both low- and high-level features that help improve music memorability. These quantitative findings manifest that data-driven MIR tasks are notably reliant on huge data quantities to be resilient and general.

5.2 Ablation Study

Table 4 shows an ablation study on feature selection for handcrafted features, indicating that selecting top-25 features leads to the best overall correlation results. Moreover, Table 3 also shows an ablation study on extra pitch shifting for data augmentations. Small pitch shifts (less than 5 semitones) make the altered audio seem natural to the human ear according to [40]. Therefore, we add semitone shifts of -5 to 5 to our data. The results manifest that pitch shifting is effective for the models that take sequence information into account because applying mean pooling across time on harmony features in non-sequential models like SVR and MLP just forces the model to forecast the same value using multiple static chroma information. This may confuse the model on harmony characteristics. On the other hand, models with sequential information, such as SSAST, learn pitch-invariance after pitch shifting. The performance of SSAST notably decreases without pitch shift data augmentation, possibly due to its data-hungry nature as a Transformer-based model, *i.e.* requiring more data for optimal parameter tuning.

5.3 Interpretability

We attempt to gain insight into the intrinsic memory utilizing XAI methodologies. One post-hoc strategy for expressing black box models in a human-interpretable manner is SHAP [16]. Specifically, SHAP explanations are obtained by perturbing a specific instance in the data and observing the impact of these perturbations on the black-box




Figure 5. SHAP summary of the SVR model with RBF kernel [41]. The most important features are listed in decreasing order and the fact that feature value rises after the SHAP value shows a positive relationship between the two.

model’s output. As such, SHAP allows us to explore the factors that the model considers when determining memorability. Figure 5 visualizes the directionality impact of the top-5 features in SHAP, where the x-axis stands for SHAP value and each point is a SHAP value of a sample for a feature. Red color and blue colors respectively indicate higher and lower values of a feature. As such, we can observe the feature directionality impact based on the distribution. For example, the first row shows that a higher arousal value leads to high memorability scores, while a lower arousal value can lead to both high and low memorability scores. The important factors for the predictor among the EHC features include arousal, bpm, harmony (the feature "D mean") and the timbre features extracted from the source other than vocals, drums, and bass (the feature "other db mean"). According to Psychology research [28], normal individuals without brain damage find it easier to recognize musical excerpts with high arousal. The melodies are the main constituent elements of the source "others" after applying 4-stem Spleeter separation. This finding supports our understanding that we often focus on the main melody in music, and thus the chorus or hook of the song with outstanding melody usually represents the entire song.

6. CONCLUSION AND FUTURE WORK

In this work, we explore the novel task of music memorability regression (MMR) using a data-driven approach. The consistency of our newly proposed YouTube Music Memorability (YTMM) dataset supports our hypothesis that music memorability indeed exists and can be predicted. Furthermore, we investigate the use of feature engineering and self-supervised learning for predicting music memorability, highlighting the importance of prior knowledge and other training approaches, such as label normalization, for improving results with limited data. We make the dataset available online to encourage further research and development in the field of MMR. In the future, we plan to: 1) scale the dataset to better represent the memorability of full music structures, 2) investigate the potential of transfer learning trained on music-oriented datasets to enhance our current baselines, and 3) study the personalization issue since music memorability can be strongly related to the past musical experience of individuals.

7. ACKNOWLEDGEMENT

This work was supported in part by the National Science and Technology Council of Taiwan under Grant NSTC-109-2221-E-009-114-MY3.

8. REFERENCES

- [1] M. Alexomanolaki, C. Loveday, and C. Kennett, "Music and memory in advertising: Music as a device of implicit learning and recall," *Music, Sound, and the Moving Image*, vol. 1, no. 1, pp. 51–71, 2007.
- [2] S. Hecker, "Music for advertising effect," *Psychology & Marketing*, vol. 1, no. 3-4, pp. 3–8, 1984.
- [3] B. Snyder and R. Snyder, *Music and memory: An introduction*. MIT press, 2000.
- [4] D. B. Ramsay, I. Ananthabhotla, and J. A. Paradiso, "The intrinsic memorability of everyday sounds," *AES International Conference on Immersive and Interactive Audio*, 2019.
- [5] A. R. Halpern and D. Müllensiefen, "Effects of timbre and tempo change on memory for music," *Quarterly Journal of Experimental Psychology*, vol. 61, no. 9, pp. 1371–1384, 2008.
- [6] D. Müllensiefen and A. R. Halpern, "The role of features and context in recognition of novel melodies," *Music Perception: An Interdisciplinary Journal*, vol. 31, no. 5, pp. 418–435, 2012.
- [7] S. McCullough Campbell and E. H. Margulis, "Catching an earworm through movement," *Journal of New Music Research*, vol. 44, no. 4, pp. 347–358, 2015.
- [8] K. Jakubowski, S. Finkel, L. Stewart, and D. Müllensiefen, "Dissecting an earworm: Melodic features and song popularity predict involuntary musical imagery." *Psychology of Aesthetics, Creativity, and the Arts*, vol. 11, no. 2, p. 122, 2017.
- [9] P. Isola, J. Xiao, D. Parikh, A. Torralba, and A. Oliva, "What makes a photograph memorable?" *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 7, pp. 1469–1482, 2013.
- [10] A. Khosla, A. S. Raju, A. Torralba, and A. Oliva, "Understanding and predicting image memorability at a large scale," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [11] R. Dubey, J. Peterson, A. Khosla, M.-H. Yang, and B. Ghanem, "What makes an object memorable?" in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1089–1097.
- [12] S. Shekhar, D. Singal, H. Singh, M. Kedia, and A. Shetty, "Show and recall: Learning what makes videos memorable," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2017, pp. 2730–2739.
- [13] A. Adadi and M. Berrada, "Peeking inside the black-box: a survey on explainable artificial intelligence (xai)," *IEEE access*, vol. 6, pp. 52 138–52 160, 2018.
- [14] C.-Y. Li, P.-C. Yuan, and H.-Y. Lee, "What does a network layer hear? analyzing hidden representations of end-to-end asr through speech synthesis," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6434–6438.
- [15] J. V. Jeyakumar, J. Noor, Y.-H. Cheng, L. Garcia, and M. Srivastava, "How can i explain this to you? an empirical study of deep neural network explanation methods," *Advances in Neural Information Processing Systems*, vol. 33, pp. 4211–4222, 2020.
- [16] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. Curran Associates, Inc., 2017, pp. 4765–4774. [Online]. Available: <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
- [17] J. A. Burgoyne, D. Bountouridis, J. van Balen, H. Honing *et al.*, "Hooked: a game for discovering what makes music catchy," in *Proceedings of the 14th Society of Music Information Retrieval Conference (ISMIR)*, 2013.
- [18] J. V. Balen, J. A. Burgoyne, D. Bountouridis, D. Müllensiefen, and R. C. Veltkamp, "Corpus analysis tools for computational hook discovery," in *ISMIR*, 2015, pp. 227–233.
- [19] J. Van Balen *et al.*, "Audio description and corpus analysis of popular music," Ph.D. dissertation, Utrecht University, 2016.
- [20] I. R. Korsmit, J. A. Burgoyne, and H. Honing, "If you wanna be my lover... a hook discovery game to uncover individual differences in long-term musical memory," in *Proceedings of the 25th Anniversary Conference of the European Society for the Cognitive Sciences of Music*, 2017.
- [21] B. McFee, E. J. Humphrey, and J. P. Bello, "A software framework for musical data augmentation." in *ISMIR*, vol. 2015, 2015, pp. 248–254.
- [22] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, "Randaugment: Practical automated data augmentation with a reduced search space," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 702–703.
- [23] H.-H. Wu, C.-C. Kao, Q. Tang, M. Sun, B. McFee, J. P. Bello, and C. Wang, "Multi-task self-supervised pre-training for music classification," in *ICASSP 2021- IEEE International Conference on Acoustics,*

- Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 556–560.
- [24] R. Castellon, C. Donahue, and P. Liang, “Codified audio language modeling learns useful representations for music information retrieval,” in *ISMIR*, 2021.
- [25] S. Pascual, M. Ravanelli, J. Serrà, A. Bonafonte, and Y. Bengio, “Learning Problem-Agnostic Speech Representations from Multiple Self-Supervised Tasks,” in *Proc. of the Conf. of the Int. Speech Communication Association (INTER-SPEECH)*, 2019, pp. 161–165. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-2605>
- [26] M. Ravanelli, J. Zhong, S. Pascual, P. Swietojanski, J. Monteiro, J. Trmal, and Y. Bengio, “Multi-task self-supervised learning for robust speech recognition,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6989–6993.
- [27] C. Spearman, “The proof and measurement of association between two things,” *The American journal of psychology*, vol. 100, no. 3/4, pp. 441–471, 1987.
- [28] S. Samson, D. Dellacherie, and H. Platel, “Emotional power of music in patients with memory disorders: Clinical implications of cognitive neuroscience,” *Annals of the New York Academy of Sciences*, vol. 1169, no. 1, pp. 245–255, 2009.
- [29] P. Vuilleumier and W. Trost, “Music and emotions: from enchantment to entrainment,” *Annals of the New York Academy of Sciences*, vol. 1337, no. 1, pp. 212–222, 2015.
- [30] C. Harte and M. Sandler, “Automatic chord identification using a quantised chromagram,” in *Audio Engineering Society Convention 118*. Audio Engineering Society, 2005.
- [31] R. Hennequin, A. Khlif, F. Voituret, and M. Moussallam, “Spleeter: a fast and efficient music source separation tool with pre-trained models,” *Journal of Open Source Software*, vol. 5, no. 50, p. 2154, 2020.
- [32] A. T. Cemgil, B. Kappen, P. Desain, and H. Honing, “On tempo tracking: Tempogram representation and kalman filtering,” *Journal of New Music Research*, vol. 29, no. 4, pp. 259–273, 2000.
- [33] K. Zhang, H. Zhang, S. Li, C. Yang, and L. Sun, “The pmemo dataset for music emotion recognition,” in *Proceedings of the 2018 acm on international conference on multimedia retrieval*, 2018, pp. 135–142.
- [34] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, “Panns: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [35] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [36] S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2009.
- [37] Y. Gong, C.-I. J. Lai, Y.-A. Chung, and J. Glass, “Ssast: Self-supervised audio spectrogram transformer,” in *AAAI Conference on Artificial Intelligence*, 2022.
- [38] K. Choi, G. Fazekas, M. Sandler, and K. Cho, “Transfer learning for music classification and regression tasks,” in *The 18th International Society of Music Information Retrieval (ISMIR) Conference 2017, Suzhou, China*. International Society of Music Information Retrieval, 2017.
- [39] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *ICLR (Poster)*, 2015.
- [40] J. Thickstun, Z. Harchaoui, D. P. Foster, and S. M. Kakade, “Invariances and data augmentation for supervised music transcription,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 2241–2245.
- [41] I. Steinwart, D. Hush, and C. Scovel, “An explicit description of the reproducing kernel hilbert spaces of gaussian rbf kernels,” *IEEE Transactions on Information Theory*, vol. 52, no. 10, pp. 4635–4643, 2006.