

TRACING BACK MUSIC EMOTION PREDICTIONS TO SOUND SOURCES AND INTUITIVE PERCEPTUAL QUALITIES

Shreyan CHOWDHURY (shreyan.chowdhury@jku.at)¹, Verena PRAHER (0000-0001-5466-7829)¹, and
Gerhard WIDMER (0000-0003-3531-1282)^{1,2}

¹*Institute of Computational Perception, Johannes Kepler University, Linz, Austria*

²*LIT AI Lab, Linz Institute of Technology, Linz, Austria*

ABSTRACT

Music emotion recognition is an important task in MIR (Music Information Retrieval) research. Owing to factors like the subjective nature of the task and the variation of emotional cues between musical genres, there are still significant challenges in developing reliable and generalizable models. One important step towards better models would be to understand what a model is actually learning from the data and how the prediction for a particular input is made. In previous work, we have shown how to derive explanations of model predictions in terms of spectrogram image segments that connect to the high-level emotion prediction via a layer of easily interpretable perceptual features. However, that scheme lacks intuitive musical comprehensibility at the spectrogram level. In the present work, we bridge this gap by merging audioLIME – a source-separation based explainer – with mid-level perceptual features, thus forming an intuitive connection chain between the input audio and the output emotion predictions. We demonstrate the usefulness of this method by applying it to debug a biased emotion prediction model.

1. INTRODUCTION

The quest for interpreting the inner workings of “black-box” models and explaining their predictions has led to many recent advances in the area of explainable AI and is becoming an increasingly important staple of all AI subfields. Not only do model explanations help in enhancing trust in the model in applications where its predictions are critical decisions like creditworthiness or medical diagnosis, but they can also often reveal telling signs of algorithmic bias [1, 2]. While algorithmic decisions in the field of MIR are not as life-critical as medical diagnosis, in today’s era of music streaming and recommendations, they can have far-reaching effects on diverse audiences, creators, and artists alike. For instance, machine-learning-based music recommender systems have been shown to exhibit severe biases towards or against certain user groups [3]. Interpretable explanations of these models or their predictions would be extremely helpful for identifying such

biases.

Previous work on interpretability in MIR has dealt with tasks such as music tagging using self-attention [4] and transcription using invertible neural networks [5], and post-hoc explanations for music content analysis have been used to understand what a genre classifier [6] or a singing voice detector [7–11] have learnt. More recently, audioLIME has been proposed [12, 13] and has shown promise in explaining tagging models [14] as well as recommendation models [15].

However, explanations of *music emotion recognition* systems have received relatively less attention notwithstanding the importance of this task in the areas of musical analysis and recommendation. [16] used models trained with different combinations of sound sources to deduce the importance of each source to emotion predictions. We previously proposed using an intermediate layer of *mid-level perceptual features* [17] to explain music emotion recognition models through a linear connection between the intermediate and final layers [18]. We followed this up with a two-step explanation approach [19] to further explain the predictions in the mid-level layer using components from the input spectrogram. This two-level scheme used LIME (*Local Interpretable Model-agnostic Explanations*) [13] to construct explanations for the mid-level predictions in terms of specific patches of the input spectrogram (which were obtained via image segmentation). While this gave us the regions in the spectrogram contributing most to a particular prediction, these regions did not hold any musical meaning by themselves. As a result, it is difficult to comprehend these explanations in terms of meaningful or intuitive concepts.

In this work, we bridge this gap by merging the audioLIME method, which uses sound sources as explanatory features, with the approach of mid-level features, to obtain comprehensible explanations from the input audio as well as from the perceptual layer. It thus forms an intuitive connection of hierarchical explanations from low-level constituent sources of audio to the high-level emotion predictions through the intermediary mid-level layer, all of which have a musical interpretation.

We believe that explainability is particularly important for developing better music emotion recognition algorithms since it is often difficult to identify misclassifications and biases in this task because of its inherent subjectivity and inter-annotator variability. As an example of real-life application of our method in understanding the po-

Copyright: © 2021 the Authors. This is an open-access article distributed under the terms of the [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

tential cause of bias in an emotion model, we demonstrate how a model that has seen few examples of a genre during training results in a pattern of errors on the test set that contains examples of this genre. We trace this pattern of errors back to a particular mid-level feature and this feature to a particular source in the input. Doing so allows us to predict how the model would change when retrained with a balanced training set. We can then qualitatively verify that the retrained model has in fact changed in the way that we expected from our explanations.

2. TWO-LEVEL EXPLANATIONS

Our proposed two-level system will explain emotion predictions by first tracing them back to the most relevant mid-level features and in a second step explain the intermediate mid-level layer via audio sources. We will first describe each of the parts separately from the lowest level (audio sources) to the highest level (emotion predictions) and then put together all the parts in Section 2.3.

2.1 Explaining via Audio Sources: audioLIME

In order to explain individual mid-level predictions we make use of audioLIME, a recently introduced approach based on LIME [13] for interpreting models in MIR [12, 14]. LIME uses simplified inputs based on a set of human interpretable features (depending on the domain of the task – e.g., superpixels for images) to train a simpler explanation model g in order to explain a more complex, potentially deep model f . Previous approaches based on LIME have used time-, frequency-, or time-frequency segments [7], or segments computed by an image segmentation algorithm [19]. audioLIME introduced a new type of interpretable features: sound sources estimated by a music source separation algorithm. In other words, the audioLIME explanation for a given prediction will be in the form of particular sound sources (and possibly specific temporal segments – which we do not use right now), telling us that it is some sonic aspects of these sources that seem to be influential. In our case, the source separator is the pretrained music source separator spleeter [20], thus the explanatory sound sources will be (what the source separator believes are) individual instruments.

2.2 Explaining via Mid-level Perceptual Features

Mid-level features are perceptual qualities or descriptors that emerge from low-level musical building blocks such as timbre, beat structure, harmony, etc. They are more subjective than the low-level features, thus difficult to model using hand-crafted feature extractors, but musically discernible enough to have many people reach a high agreement in annotations. Examples include qualities such as perceived melodiousness or rhythmic complexity [17]. They, therefore, form a suitable choice for an intermediary to higher-level concepts like emotion.

This idea was first used in [18] to explain emotion predictions. We use a similar approach here, but we use the more recently introduced receptive-field regularized ResNets [21, 22] to model the emotions and the mid-level

features. In addition, we learn the emotions and mid-level features from two separate datasets using multi-task learning, i.e., jointly learn to predict mid-level features and high-level emotions, allowing us to be flexible with our train and test domains.

The mid-level layer is the penultimate layer of the model, with a linear transformation between it and the final (emotion) layer, thus making the connection interpretable. It is learnt end-to-end from audio spectrograms by optimizing on the combined loss from the emotion and mid-level layers.

Emotion predictions can be interpreted by looking at *effects plots*, which are visual representations of the contribution of each mid-level feature to the final emotion prediction, calculated as the product of the feature value and the weight joining it to the emotion. These indicate the relative influence of the individual features on the final prediction.

2.3 Putting it All Together: Intuitive Two-level Explanations

For a given prediction, we then work backwards. First, we obtain the mid-level explanation of an emotion by computing the effects. The larger the effect of a mid-level feature, the larger is the contribution of that feature to the emotion prediction. Next, we compute the audioLIME explanations for the mid-level feature with the largest effect (we can in principle compute audioLIME explanations for all features to obtain a more diverse explanation, depending on the application). Given these two explanations, we can describe a prediction as being arrived at by the model due to the explanatory mid-level feature, which is in turn most influenced by the input component given by audioLIME.

3. EXPERIMENTAL SETUP

For our analysis with emotion explanations, we first need to train the “explainable” models, which have the penultimate mid-level layer connecting linearly to the final emotion outputs. This section describes the datasets and the training procedure for such models.

3.1 Datasets

For our experiments we are using three different datasets:

3.1.1 Mid-level Perceptual Features Dataset

The *Mid-level Perceptual Features Dataset* introduced in [17] consists of 5000 song snippets with annotations between 1 and 10 for the mid-level descriptors *melodiousness*, *articulation*, *rhythmic complexity*, *rhythmic stability*, *dissonance*, *tonal stability*, and *modality* (called “minor-ness” here). We use this dataset to train the intermediate layer of our emotion model.

3.1.2 DEAM: Database for Emotional Analysis in Music

The *DEAM* dataset [23] is a dataset of dynamic and static valence and arousal annotations. It contains 1,802 songs (58 full-length songs and 1,744 excerpts of 45 seconds) from a variety of Western popular music genres (rock, pop,

	Arousal		Valence	
	RMSE	R2	RMSE	R2
P(P - bl)	0.23	0.61	0.25	0.41
P(P)	0.25 ± 0.03	0.60 ± 0.10	0.31 ± 0.04	0.40 ± 0.14
D(P)	0.27 ± 0.01	0.50 ± 0.03	0.33 ± 0.00	0.30 ± 0.02
D(D)	0.26 ± 0.01	0.49 ± 0.02	0.22 ± 0.01	0.51 ± 0.04
D+P(P)	0.23 ± 0.01	0.65 ± 0.02	0.28 ± 0.00	0.50 ± 0.02
D+P(D)	0.26 ± 0.01	0.50 ± 0.03	0.23 ± 0.01	0.48 ± 0.02

Table 1: Emotion prediction performance with our “explainable” model trained and tested on difference datasets – P: PMEmo, D: DEAM. The dataset inside the parentheses is the test dataset. The top row is the baseline performance from [16].

electronic, country, jazz, etc). In our experiments, we use the static emotion annotations, which are continuous values between 0 and 10.

3.1.3 PMEmo: Popular Music with Emotional Annotation

The PMEmo dataset [24] consists of 794 chorus clips from three different well-known music charts. The songs were annotated by 457 annotators with valence and arousal annotations separately for dynamic and static. In our experiments, we use static labels, which are continuous values between 0 and 1.

3.2 Model Training

The mid-level and emotion model is trained end-to-end using audio spectrograms as inputs and optimizing on the combined loss from the mid-level and emotion layers. The batch size is 16 and contains 8 samples from the Mid-level dataset and 8 samples from either the DEAM or the PMEmo dataset. The loss function is the mean squared error. The learning rate is 10^{-3} with cosine annealing, and we perform early stopping on a validation set as regularization. We use the Adam optimizer [25].

The inputs are log-filtered spectrograms (149 bands) of 40-second audio clips peak normalized and sampled at 22.05 kHz with a window size of 2048 samples and a hop length of 704 samples, resulting in 149×1252 -sized tensors. If a clip is longer than 40 seconds, we take a random snippet, and if it is shorter, it is looped to 40 seconds.

The labels are scaled to the range $[-1, 1]$ for all three datasets. Therefore, an RMSE of 0.26 would represent 13% error. We split the train and test sets such that they have mutually exclusive sets of artists. A summary of the emotion prediction performance can be found in Table 1.

4. EVALUATION OF EXPLANATIONS

Essentially, there are three targets for empirical evaluation: the two individual components of our two-level explanation framework, and the final composite explanations produced by the model. Regarding the former, the higher level – explanations of emotion predictions in terms of mid-level perceptual features – has already been discussed at length in our previous paper [18]. We showed how effects plots can give insight into the relative importance of various mid-level qualities. The lower level – using audioLIME to explain mid-level feature predictions via audio sources – is

a new concept, and the experiments in the following section are intended to validate it. Empirical evidence for the usefulness of the complete, two-level explanation model, finally, will be presented in the form of a study, in Section 5 below, where we demonstrate how these explanations can help us debug a biased prediction model by gaining insight into what the sources of its problems are.

4.1 Explaining Mid-level Features via Sound Sources

Evaluating the quality of explanations is a hard task since there is no consensus on what makes a good explanation, with a variety of desired aims and properties proposed in literature [26–28]). We build our evaluation of audioLIME explanations for the mid-level layer on two metrics, (a) *fidelity* as proposed by Ribeiro et al. along with LIME [13] and, (b) *complexity*, a recently proposed metric for feature-based model explanations [26].

Fidelity measures how well the local model g (the explainer) approximates the global model f (the model up to the mid-level layer in our case) [13] and is computed using the coefficient of determination between the local and global model’s predictions as in the original LIME implementation¹.

In addition to high fidelity, low complexity is desired. The most complex explanation would be the one where all d features get the same attribution (i.e., all weights $g(f, x)_i$ of the linear explanation model g are the same). The simplest explanation concentrates all attribution on one feature. To measure complexity, a probability distribution P_g is defined:

$$P_g(i) = \frac{|g(f, x)_i|}{\sum_{j \in [d]} |g(f, x)_j|} \quad (1)$$

Complexity is then defined as the entropy of this distribution [26]. We compare the complexity per dataset with a random baseline, which is obtained by creating “random” explanations with feature weights drawn from a uniform distribution.

For the analysis, we compute predictions and explanations for all test examples and calculate the above mentioned metrics. The results for one mid-level feature (we picked “rhythmic stability” as it is used later on as an example) are summarized in Figures 1 and 2. We can see in Figure 1a that the fidelity score (coefficient of determination) is relatively high across all combinations of models and test sets. The median score is 0.86 across all explanations (including all mid-level features), the 25%-quantile is at 0.78. This means that for 50% and 75% of the explanations more than 86%, and 78%, respectively, of the variation in the dependent variable (mid-level prediction) can be predicted using the independent variables (instrument sources).

Figure 1b shows the computed complexities, compared to a random baseline. Most explanations are far less complex than the random baseline.

The results shown in the previous figures suggested a relationship between the fidelity and complexity scores.

¹ <https://github.com/marcotcr/lime/>

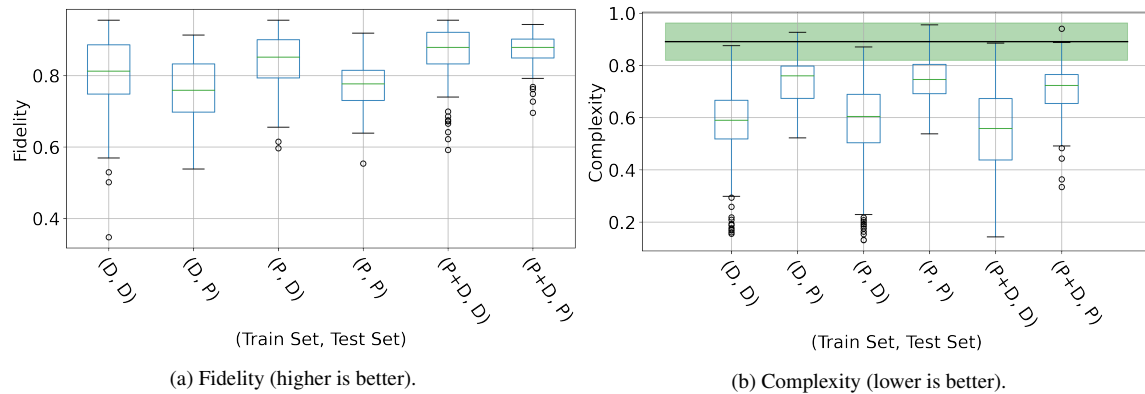


Figure 1: Figure 1a shows the computed fidelity (coefficient of determination R^2 between the predictions by the global model f and the local model g) scores for the evaluated explanations. Figure 1b shows the complexity (entropy of a distribution over the feature attribution weights) scores for the evaluated explanations. The green region shows the standard deviation of complexities for 1000 random explanations, with the black line being the mean.

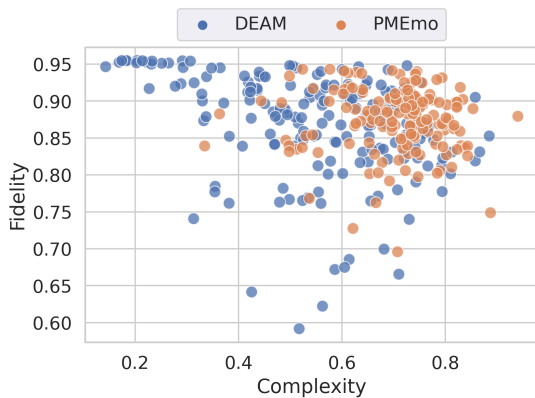


Figure 2: A more detailed view on the relationship between the fidelity score and complexity for the predictions of “rhythmic stability” for a model trained on both datasets. The color indicates the test set.

Therefore we visualized the two metrics for all explanations computed for “rhythmic stability” for a model trained on the combined data sets in Figure 2. Although they seem related on a dataset level, the metrics do not look related when analyzed for each explanation separately, suggesting that indeed both are needed.

5. MODEL DEBUGGING

A practical use case of our explanation scheme is demonstrated in this section. We use the two-level explanations to understand why an improperly trained model might be overestimating the valence predictions for one particular genre.

5.1 Setup

First, we use a pre-trained tagger [29] to predict genre tags for all the tracks in the three datasets mentioned in Section 3.1, since we do not have genre metadata for these

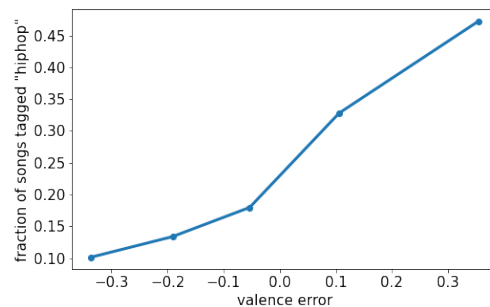


Figure 3: Fraction of hip-hop songs in quantiles vs the mean valence error of each quantile over PMemo dataset (with model trained on DEAM)

datasets. This is only done in order to obtain an estimate of the genre-dependence of the emotion predictions later on. We then train two explainable models – one on the DEAM dataset, and one on the combined DEAM and PMemo dataset. The test set is a fixed but randomly chosen subset of the PMemo dataset (with a mutually exclusive set of artists from the training set).

5.2 Overestimated Valence for Hip-hop

When we take the model trained only on DEAM and use it to predict arousal and valence for the entire PMemo dataset, we observe that the error in valence shows a pattern – overestimations of valence primarily occur in hip-hop songs, as shown in Figure 3.

We can reason about relatively poor performance for hip-hop songs based on the discrepancy between the training and testing sets in terms of genre composition. In Figure 5, we can see that PMemo has a large percentage of hip-hop songs whereas both DEAM and Mid-level datasets have a small percentage. Since our model has not seen enough hip-hop songs during training, it is to be expected that it does not perform well when it encounters hip-hop

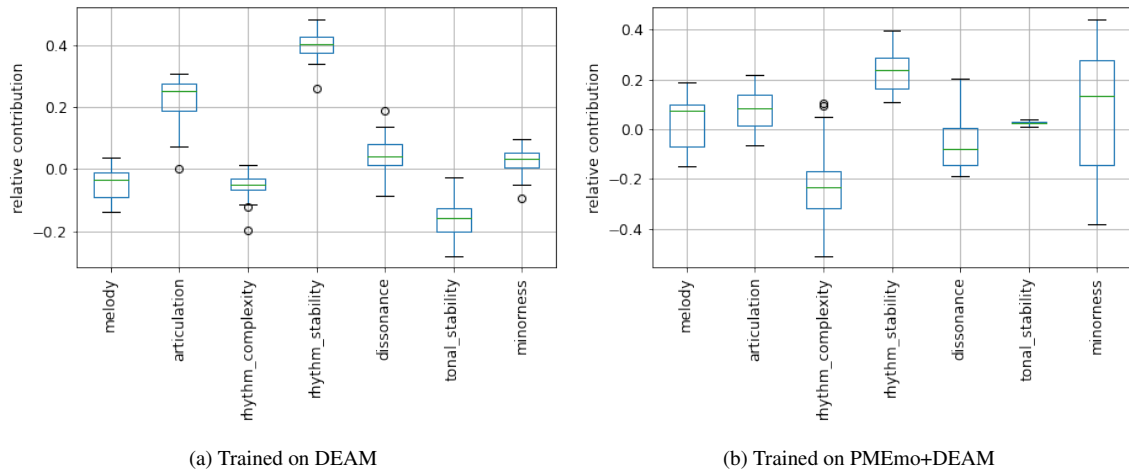


Figure 4: Relative effects of the mid-level features for valence prediction for two models trained on different datasets, but tested on the same fixed subset of the PMEmo dataset.

during test. However, a question that is pertinent next – what is it about hiphop songs that makes our model overestimate their valence?

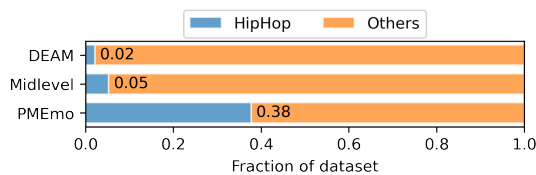


Figure 5: Compositions of datasets as fraction of songs tagged “hiphop” by a pre-trained auto-tagging model [29]

5.3 Explaining Valence Overestimations Using Mid-level Features

To answer this question, we first seek to understand which of the mid-level qualities can be attributed most to high valence predictions. This is the first level of our explanation system. We find these attributions by computing the effects of each mid-level feature on the valence predictions. The effect of a feature is simply the value of that feature multiplied by the weight of the linear connection between it and the target node. In our case, the target is valence and there are seven mid-level features that affect it. We are only interested in relative contribution of each feature, and so we divide each effect by the sum of the absolute values of the effects of all features and take the average across all test songs tagged “hiphop”.

We observe that rhythmic stability has the maximum positive relative effect on the prediction of valence. Therefore, we select rhythmic stability for the next step of explanation.

5.4 Explaining Rhythmic Stability Using Sources

Once we have selected a mid-level feature as having the most positive relative effect on the valence, we would like

to understand what musical constituents in the input can be attributed to positive contribution to that feature. To do this, we take the help of audioLIME and generate source based explanations for rhythmic stability. The sources available in the current implementation of audioLIME² are vocals, drums, bass, piano, and other.

We find that vocals are a major contributing source for the rhythmic stability predictions for the hiphop songs. For songs tagged as other genres, contributing sources are more distributed.

5.5 Re-training the Model with Target Data

Bringing together our two types of explanations, we can reason that the high valence predictions for hiphop songs is due to overestimation of rhythmic stability, which, in this case, can be attributed to the vocals. While there is a lot of diversity in the style of *rapping* (the form of vocal delivery predominant in hiphop), it has been noted that rappers typically use stressed syllables and vocal onsets to match the vocals with the underlying rhythmic pulse [30,31]. These rhythmic characteristics of vocal delivery (that constitutes “flow”, and may add metrical layers on top of the beat) contribute strongly to the rhythmic feel of a song. The positive or negative emotion of hiphop songs is mostly contained in the lyrics – the style of vocal performance does not necessarily express or correlate with this aspect of emotion. Therefore, it makes sense that a model which has seen few examples of hiphop during training should wrongly associate the prominent rhythmic vocals of hiphop to high rhythmic stability and in turn high valence. A model that has been trained with hiphop songs included, we expect, would place less importance on rhythmic stability for the prediction of valence, even if the vocals might still contribute significantly to rhythmic stability. Thus, we expect the relative effect of rhythmic stability for valence to decrease in such a model.

² <https://github.com/CPJKU/audioLIME>

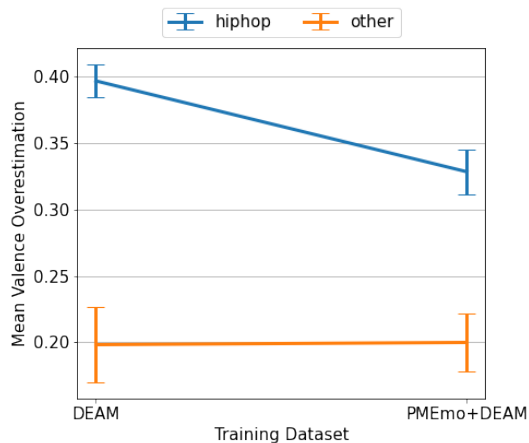


Figure 6: Mean valence overestimations for two models trained on different datasets, but tested on the same fixed subset of the PMEmo dataset.

This is exactly what we observe on a model trained with the combined PMEmo+DEAM dataset. The average relative effects are shown in Figure 4b and we can see that the relative effect of rhythmic stability has decreased while those of minoriness, melody, and tonal stability have increased. Thus, the model changed in a way that was in line with what we expected from the analysis of our two-level explanation method.

Looking at mean overestimations (Figure 6) in valence for hip-hop and other genres for models trained on DEAM and PMEmo+DEAM shows that valence overestimations of hip-hop songs have decreased substantially, without negatively affecting the predictions on other genres³.

6. CONCLUSION AND FUTURE WORK

In this paper, we proposed a method to explain music emotion models in an intuitive way using components from low- and mid- levels of the hierarchy of musical concepts by combining audioLIME, which uses input audio sources as explanatory components, with intermediate layer based explanations. We also demonstrated its potential as a tool for model debugging and explaining model behaviour.

This points us towards exploring this method further and getting more granular explanations as a way of improving the effectiveness of this system for MIR. An immediate next step that we are currently pursuing is to extend audioLIME to provide explanations in the form of temporal segments using semantic music segmentation, along with the sound sources.

We are also looking at explaining emotion conveyed in classical piano performances, which pose particular challenges – including the non-availability of training data, where transfer learning of explanatory features becomes necessary [32].

³ Code for reproducing model debugging experiments is available at https://github.com/shreyanc/model_debugging

Acknowledgments

This work is supported by the European Research Council (ERC) under the EU’s Horizon 2020 research & innovation programme under grant agreement No. 670035 (“Con Espressione”), and the Federal State of Upper Austria (LIT AI Lab).

7. REFERENCES

- [1] L. Rieger, C. Singh, W. J. Murdoch, and B. Yu, “Interpretations are Useful: Penalizing Explanations to Align Neural Networks with Prior Knowledge,” in *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, vol. 119. PMLR, 2020, pp. 8116–8126.
- [2] W. Samek and K. Müller, “Towards Explainable Artificial Intelligence,” in *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, ser. Lecture Notes in Computer Science. Springer, 2019, vol. 11700, pp. 5–22.
- [3] M. D. Ekstrand, M. Tian, I. M. Azpiazu, J. D. Ekstrand, O. Anuyah, D. McNeill, and M. S. Pera, “All The Cool Kids, How Do They Fit In?: Popularity and Demographic Biases in Recommender Evaluation and Effectiveness,” in *Conference on Fairness, Accountability and Transparency, FAT 2018, 23-24 February 2018, New York, NY, USA*, vol. 81. PMLR, 2018, pp. 172–186.
- [4] M. Won, S. Chun, and X. Serra, “Toward Interpretable Music Tagging with Self-Attention,” *CoRR*, vol. abs/1906.04972, 2019.
- [5] R. Kelz and G. Widmer, “Towards Interpretable Polyphonic Transcription with Invertible Neural Networks,” in *Proceedings of the 20th International Society for Music Information Retrieval Conference, ISMIR 2019, Delft, The Netherlands, November 4-8, 2019*, 2019, pp. 376–383.
- [6] K. Choi, G. Fazekas, and M. B. Sandler, “Explaining Deep Convolutional Neural Networks on Music Classification,” *CoRR*, vol. abs/1607.02444, 2016.
- [7] S. Mishra, B. L. Sturm, and S. Dixon, “Local Interpretable Model-Agnostic Explanations for Music Content Analysis,” in *Proceedings of the 18th International Society for Music Information Retrieval Conference, ISMIR 2017, Suzhou, China, October 23-27, 2017*, 2017, pp. 537–543.
- [8] S. Mishra, E. Benetos, B. L. Sturm, and S. Dixon, “Reliable Local Explanations for Machine Listening,” in *2020 International Joint Conference on Neural Networks, IJCNN 2020, Glasgow, United Kingdom, July 19-24, 2020*. IEEE, 2020, pp. 1–8.
- [9] S. Mishra, B. L. Sturm, and S. Dixon, “Understanding a Deep Machine Listening Model Through Feature Inversion,” in *Proceedings of the 19th International Society for Music Information Retrieval Conference, ISMIR*

- 2018, Paris, France, September 23-27, 2018, 2018, pp. 755–762.
- [10] S. Mishra, D. Stoller, E. Benetos, B. L. Sturm, and S. Dixon, “GAN-based Generation and Automatic Selection of Explanations for Neural Networks,” *CoRR*, vol. abs/1904.09533, 2019.
- [11] S. Mishra, B. L. Sturm, and S. Dixon, ““What are You Listening to?” Explaining Predictions of Deep Machine Listening Systems,” in *26th European Signal Processing Conference, EUSIPCO 2018, Roma, Italy, September 3-7, 2018*. IEEE, 2018, pp. 2260–2264.
- [12] V. Haunschmid, E. Manilow, and G. Widmer, “audioLIME: Listenable Explanations Using Source Separation,” in *13th International Workshop on Machine Learning and Music*, 2020, pp. 20–24.
- [13] M. T. Ribeiro, S. Singh, and C. Guestrin, ““Why Should I Trust You?”: Explaining the Predictions of Any Classifier,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*. ACM, 2016, pp. 1135–1144.
- [14] V. Haunschmid, E. Manilow, and G. Widmer, “Towards Musically Meaningful Explanations Using Source Separation,” *CoRR*, vol. abs/2009.02051, 2020.
- [15] A. B. Melchiorre, V. Haunschmid, M. Schedl, and G. Widmer, “LEMONS : Listenable Explanations for Music recOmmeNder Systems,” in *Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR 2021 (forthcoming), Lucca, Italy, March 28 – April 1, 2021*.
- [16] J. de Berardinis, A. Cangelosi, and E. Coutinho, “The Multiple Voices of Musical Emotions: Source Separation for Improving Music Emotion Recognition Models and Their Interpretability,” *Proceedings of the 21st International Society for Music Information Retrieval Conference*, no. iii, pp. 310–217, 2020.
- [17] A. Aljanaki and M. Soleymani, “A Data-driven Approach to Mid-level Perceptual Musical Feature Modeling,” in *Proceedings of the 19th International Society for Music Information Retrieval Conference, ISMIR 2018, Paris, France, 2018*, pp. 615–621.
- [18] S. Chowdhury, A. Vall, V. Haunschmid, and G. Widmer, “Towards Explainable Music Emotion Recognition: The Route via Mid-level Features,” in *Proceedings of the 20th International Society for Music Information Retrieval Conference, ISMIR 2019*, 2019.
- [19] V. Haunschmid, S. Chowdhury, and G. Widmer, “Two-level Explanations in Music Emotion Recognition,” *Machine Learning for Music Discovery Workshop, MLAMD at ICML2019*, 2019.
- [20] R. Hennequin, A. Khelif, 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, deezer Research.
- [21] K. Koutini, S. Chowdhury, V. Haunschmid, H. Eghbal-Zadeh, and G. Widmer, “Emotion and Theme Recognition in Music with Frequency-Aware RF-Regularized CNNs,” in *MediaEval 2019 Workshop, Sophia Antipolis, France, 27-30 October 2019*, vol. 2670.
- [22] K. Koutini, H. Eghbal-Zadeh, V. Haunschmid, P. Primus, S. Chowdhury, and G. Widmer, “Receptive-Field Regularized CNNs for Music Classification and Tagging,” *CoRR*, vol. abs/2007.13503, 2020.
- [23] A. Aljanaki, Y.-H. Yang, and M. Soleymani, “Developing a Benchmark for Emotional Analysis of Music,” *PloS one*, vol. 12, no. 3, p. e0173392, 2017.
- [24] 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*, ser. ICMR ’18. New York, NY, USA: ACM, 2018, pp. 135–142.
- [25] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [26] U. Bhatt, A. Weller, and J. M. F. Moura, “Evaluating and Aggregating Feature-based Model Explanations,” in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, C. Bessiere, Ed. ijcai.org, 2020, pp. 3016–3022.
- [27] Z. C. Lipton, “The Mythos of Model Interpretability,” *ACM Queue*, vol. 16, no. 3, p. 30, 2018.
- [28] N. Tintarev and J. Masthoff, “Evaluating the Effectiveness of Explanations for Recommender Systems - Methodological Issues and Empirical Studies on the Impact of Personalization,” *User Model. User Adapt. Interact.*, vol. 22, no. 4-5, pp. 399–439, 2012.
- [29] J. Pons, O. Nieto, M. Prockup, E. M. Schmidt, A. F. Ehmann, and X. Serra, “End-to-end Learning for Music Audio Tagging at Scale,” in *19th International Society for Music Information Retrieval Conference (ISMIR2018)*, Paris, 2018.
- [30] M. Ohriner, “Lyric, rhythm, and non-alignment in the second verse of Kendrick Lamar’s “Momma,”” *Music Theory Online*, vol. 25, no. 1, 2019.
- [31] K. Adams, “On the metrical techniques of flow in rap music,” *Music Theory Online*, vol. 15, no. 5, 2009.
- [32] S. Chowdhury and G. Widmer, “Towards explaining expressive qualities in piano recordings: Transfer of explanatory features via acoustic domain adaptation,” *arXiv preprint arXiv:2102.13479*, 2021.