

EMOPIA: A MULTI-MODAL POP PIANO DATASET FOR EMOTION RECOGNITION AND EMOTION-BASED MUSIC GENERATION

Hsiao-Tzu Hung^{1,2} Joann Ching¹ Seunghoon Doh³
Nabin Kim⁴ Juhan Nam³ Yi-Hsuan Yang¹

¹ Academia Sinica, Taiwan

² Department of CSIE, National Taiwan University

³ Graduate School of Culture Technology, KAIST, South Korea

⁴ Georgia Institute of Technology, United States

r08922a20@csie.ntu.edu.tw, joann8512@citi.sinica.edu, seunghoordoh@kaist.ac.kr

ABSTRACT

While there are many music datasets with emotion labels in the literature, they cannot be used for research on symbolic-domain music analysis or generation, as there are usually audio files only. In this paper, we present the EMOPIA (pronounced ‘yee-mò-pi-uh’) dataset, a shared multi-modal (audio and MIDI) database focusing on perceived emotion in pop piano music, to facilitate research on various tasks related to music emotion. The dataset contains 1,087 music clips from 387 songs and clip-level emotion labels annotated by four dedicated annotators. Since the clips are not restricted to one clip per song, they can also be used for song-level analysis. We present the procedure for building the dataset, covering the song list curation, clip selection, and emotion annotation processes. Moreover, we prototype use cases on clip-level music emotion classification and emotion-based symbolic music generation by training and evaluating corresponding models using the dataset. The result demonstrates the potential of EMOPIA for being used in future exploration on piano emotion-related MIR tasks.

1. INTRODUCTION

The affective aspect of music has been a major subject of research in the field of music information retrieval (MIR), not only for music analysis and labeling [1–9], but also for music generation or editing [10–14]. Accordingly, there have been quite a few public music datasets with emotion, as listed in Table 1. These datasets are different in many ways, including the musical genres considered, data modality and data size, and the way emotion is described.

With the growing interest in symbolic-domain music analysis and generation in recent years of ISMIR [15–19], it is desirable to have an emotion-labeled symbolic music

dataset to add emotion-related elements to such research. However, among the datasets listed in Table 1, only two provide MIDI data, and they are both small in size. Moreover, the majority of the audio-only datasets contain songs of multiple genres, making it hard to apply automatic music transcription algorithms, which currently work better for piano-only music [20–23], to get MIDI-like data from the audio recordings.

To address this need, we propose a new emotion-labeled dataset comprising of three main nice properties:

- **Single-instrument.** We collect audio recordings of piano covers and creations from YouTube, with fair to high audio and musical quality, and a diverse set of playing styles and perceived emotions. Focusing on only piano music allows for the use of piano transcription algorithms [20, 21], and facilitates disentanglement of musical composition from variations in timbre, arrangement, and other confounds seen in multi-instrument music.
- **Multi-modal.** Both the audio and MIDI versions of the music pieces can be found from the Internet (see Section 3.5 for details). The MIDI files are automatically transcribed from the audio by a state-of-the-art model [21].
- **Clip-level annotation.** The audio files downloaded from YouTube are full songs. As different parts of a song may convey different emotions, the first four authors of the paper manually and carefully pick emotion-consistent short clips from each song and label the emotion of these clips using a four-class taxonomy derived from the Russell’s valence-arousal model [32]. This leads to clip-level emotion annotations for in total 1,087 clips from 387 songs (i.e., 2.78 clips per song on average), with the number of clips per emotion class fairly balanced.

Given these properties, EMOPIA has versatile use cases in MIR research. For music labeling, EMOPIA can be used for clip-level music emotion recognition or music emotion variation detection [2], in both the audio and symbolic domains. For music generation, EMOPIA can be used for emotion-conditioned piano music generation or style transfer, to create emotion-controllable new compositions, or variations of existing pieces, again in both domains.

We present details of the dataset and the way we com-



Name	Label type	Genre (or data source)	Size	Modality
Jamendo Moods [24]	adjectives	multiple genres	18,486	Audio
DEAM [25]	VA values	(from FMA [26], Jamendo, MedleyDB [27])	1,802	Audio
EMO-Soundscapes [28]	VA values	(from FMA)	1,213	Audio
CCMED-WCMED [29]	VA values	classical (both Western & Chinese)	800	Audio
emoMusic [30]	VA values	pop, rock, classical, electronic	744	Audio
EMusic [31]	VA values	experimental, 8 others	140	Audio
MOODetector [6]	adjectives	multiple genres (AllMusic)	193	Audio+MIDI
VGMIDI [10]	valence	video game	95	MIDI
EMOPIA (ours)	Russell’s 4Q	pop (piano covers)	1,078	Audio+MIDI

Table 1. Comparison of some existing public emotion-labeled music datasets and the proposed EMOPIA dataset.

pile it in Section 3, and report a computational analysis of the dataset in Section 4. Moreover, to demonstrate the potential of the new dataset, we use it to train and evaluate a few clip-level emotion classification models using both audio and MIDI data in Section 5, and emotion-conditioned symbolic music generation models in Section 6. The latter involves the use of a recurrent neural network (RNN) model proposed by Ferreira *et al.* [10], and our own modification of a Transformer-based model [33–35] that takes emotion as a conditioning signal for generation.

We release the EMOPIA dataset at a Zenodo repo.¹ Source code implementing the generation and classification models can be found on GitHub.^{2 3} Examples of generated pieces can be found at a demo webpage.⁴

2. RELATED WORK

Emotion Recognition in Symbolic Music. Symbolic music representations describe music with the note, key, tempo, structure, chords, instruments. To understand the relationship between music and emotion, various researchers have investigated machine learning approaches with handcrafted features. Grekow *et al.* [36] extract in total 63 harmony, rhythmic, dynamic features from 83 classic music MIDI files. Lin *et al.* [37] compare audio, lyric, and MIDI features for music emotion recognition, finding that MIDI features lead to higher performance in valence dimension. Panda *et al.* [6] also proposed multi-model approaches, combining audio and MIDI features for emotion recognition using a small dataset of 193 songs.

Emotion-conditioned Symbolic Music Generation. Only few work has started to address this task recently. Ferreira *et al.* [10] compile a small dataset of video game MIDI tracks with manual annotations of valence values, named VGMIDI (cf. Table 1), and use it to train a long short term memory network (LSTM) in tandem with a genetic algorithm (GA) based elite selection mechanism to generate positive or negative music. Makris *et al.* [12] approach the same task by using designated chord progression sequence in a sequence-to-sequence architecture

trained with the VGMIDI dataset. Zhao *et al.* [38] use LSTM to generate music with four different emotions. Madhok *et al.* [13] use human facial expressions as the condition to generate music. More recently, Tan & Herremans demonstrate that their FaderNets [16] can achieve arousal-conditioned symbolic music style transfer with a semi-supervised clustering method that learns the relation between high-level features and low-level representation. Their model modifies the emotion (specifically, only the arousal) of a music piece, instead of generating new pieces from scratch.

3. THE EMOPIA DATASET

3.1 Song Selection and Segmentation

EMOPIA is a collection of 387 piano solo performances of popular music segmented manually into 1,087 clips for emotion annotation. Two authors of the paper curated the song list of the piano performances by scanning through playlists on Spotify for its consistently high quality, then downloading the recordings from YouTube. A song is included when it is played by a professional conveying a clear emotion, and the recording has not been heavily engineered during post-production. The genres of songs include Japanese anime, Korean and Western pop song covers, movie soundtracks, and personal compositions.

In an effort to extend the usefulness of the dataset for future research, at the best of our ability, the songs are intentionally segmented (with the help of the Sonic Visualizer [40]) only at cadential arrivals to make it an *emotionally-consistent clip* and a valid *musical phrase* at the same time. Accordingly, EMOPIA contains information for full songs, extracted phrases, and emotion labels.

3.2 Emotion Annotation

Different emotion taxonomies have been adopted in the literature for emotion annotation, with no standard so far [2]. For EMOPIA, we consider a simple four-class taxonomy corresponding to the four quadrants of the Russell’s famous Circumplex model of affect [32], which conceptualizes emotions in a two-dimensional space defined by valence and arousal. The four classes are: HVHA (high valence high arousal); HVLA (high valence low arousal);

¹ <https://zenodo.org/record/5090631>

² <https://github.com/annahun31/EMOPIA>

³ https://github.com/SeungHeonDoh/EMOPIA_cls

⁴ <https://annahun31.github.io/EMOPIA/>

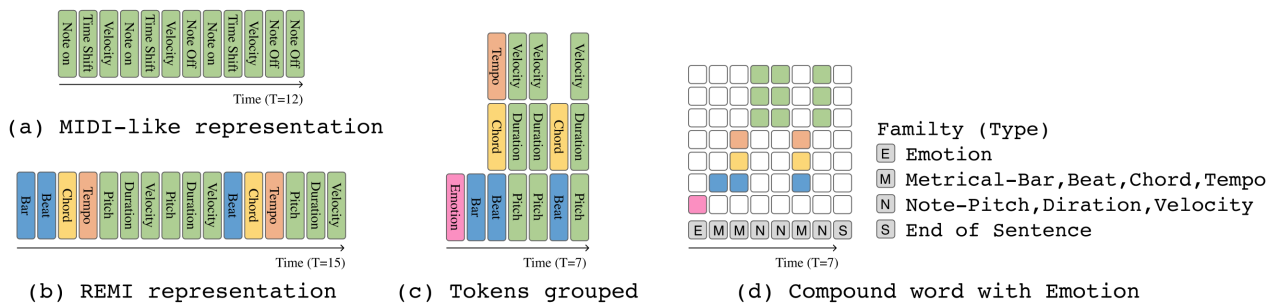


Figure 1. Illustration of different token-based representation for symbolic music: (a) MIDI-like [39], (b) REMI [34], and (d) CP [35] plus emotion token. Sub-figure (c) is an intermediate representation of the CP one.

Quadrant	# clips	Avg. length (in sec / #tokens)
Q1	250	31.9 / 1,065
Q2	265	35.6 / 1,368
Q3	253	40.6 / 771
Q4	310	38.2 / 729

Table 2. The number of clips and their average length in seconds, or in the number of REMI tokens, for each quadrant of the Russell’s model in the EMOPIA dataset.

LVHA (low valence high arousal); and LVLA (low valence low arousal). We refer to this taxonomy as Russell’s 4Q.

As pointed out in the literature [2, 3, 41], various factors affect the perceived emotion of music, including cultural background, musical training, gender, etc. Consensus on the perception of emotion is challenging accordingly. Therefore, the annotations were made only among the first four authors, all coming from similar cultural backgrounds and collaborating closely during the annotation campaign, to ensure mutual standards for high or low valence/arousal. As it is time-consuming and laborious to choose clips from a song and label the emotion, each song was taken care of by only one annotator. Yet, cross validation of the annotations among the annotators was made several times during the annotation campaign (which spans 2.5 months), to ensure all annotators work on the same standard.

Table 2 shows the number of clips and the average length (in seconds) for each class. The clips amount to approximately 11 hours’ worth of data.

3.3 Transcription

We transcribe the selected clips automatically with the help of the high-resolution piano transcription model proposed by Kong *et al.* [21], which is open source and represents the state-of-the-art for this task. We have manually checked the transcription result for a random set of clips and find the accuracy in note pitch, velocity, and duration satisfactory. The transcription might be fragmented and undesirable for cases such as when the audio recording is engineered to have unnatural ambient effects; we drop such songs from our collection. The model also transcribes pedal information, which we include to EMOPIA but do not use in our experiments.

3.4 Pre-processing and Encoding

For building machine learning models that deal with symbolic data, we need a data representation that can be used as input to the models. For example, MusicVAE [42] adopts the *event-based* representation that encodes a symbolic music piece as a sequence of “event tokens” such as note-on and note-off, while MuseGAN [43] employs a timestep-based, *piano roll*-like representation. Since there is no standard on the symbolic representation thus far, we adopt the following event-based ones in our experiments. Specifically, we use MIDI-like and REMI in Section 5 and CP in Section 6. See Figure 1 for illustrations.

- The **MIDI-like** representation [39] encodes information regarding a MIDI note with a “note-on” token, a “note-off”, and a “velocity” token. Moreover, the “time shift” token is used to indicate the relative time gap (in ms) between two tokens.
- **REMI** [34] considers a beat-based representation that instead uses “bar” and “subbeat” tokens to represent the time information. A “bar” token signifies the beginning of a new bar, and “subbeat” points to one of the constant number of subbeat divisions in a bar. Additionally, REMI uses “tempo” tokens to control the pace of the music, and replaces “note-off” with “note-duration”. Table 2 also shows the average number of REMI tokens for clips in each emotion class.
- **CP** [35]. Both MIDI-like and REMI view events as individual tokens. In CP, tokens belonging to the same *family* are grouped into a *super token* and placed on the same timestep (see Figure 1(c)). CP considers by default three families: *metrical*, *note*, and *end-of-sequence*. We additionally consider the “emotion” tokens and make it a new family, as depicted in Figure 1(d). The *prepending* approach is motivated by CTRL [44], a state-of-the-art controllable text generation model in natural language processing (NLP) that uses global tokens to affect some overall properties of a sequence.

3.5 Dataset Availability

In EMOPIA, each sample is accompanied with its corresponding metadata, segmentation annotations, emotion annotation, and transcribed MIDI, which are all available




Figure 2. Violin plots of the distribution in (a) note density, (b) length, and (c) velocity for clips from different classes.




Figure 3. Histogram of the keys (left / right: major / minor keys) for clips from different emotion classes.

Label pair	Different arousal		Different valence	
	Q1 vs. Q4	Q2 vs. Q3	Q1 vs. Q2	Q3 vs. Q4
JS Div.	0.236	0.250	0.434	0.280

Table 3. Jensen-Shannon divergence between the key histograms of a few emotion quadrant pairs.

in the Zenodo repository. Moreover, we have added the MIDI data to the MusPy library [18] to facilitate its usage. Due to copyright issues, however, we can only share audio through YouTube links instead of sharing the audio files directly; the availability of the songs are subject to the copyright licenses in different countries and whether the owners will remove them.

4. DATASET ANALYSIS

The emotion that listeners perceive is determined by a wide array of composed and performed features of music [45]. To observe the emotional correlate of the musical attributes in EMOPIA, we extract various MIDI-based features and examine the distributions over the four quadrants of emotion. We present here the most discriminative features among many choices we examined in our analysis.

Note Density, Length, and Velocity. The arousal of music can be easily observed based on the frequency of note occurrences and their strength [46]. We measure them by note density, length and velocity. The note density is defined as the number of notes per beat, and the note

length is defined as the average note length in beat unit. The note velocity is obtained directly from MIDI. Figure 2 shows three violin plots of the three features. The general trend shows that the high-arousal group (Q1, Q2) and low-arousal group (Q3, Q4) are distinguished well in all three plots. The results of note density and velocity are expected considering the nature of arousal. However, it is quite interesting that the note lengths are generally longer in the low-arousal group (Q3, Q4). Between the same-arousal quadrants, the differences are subtle. In note density, Q2 has more dynamics than Q1, whereas Q3 is not distinguishable from Q4. In note length, Q1 has slightly longer notes than those of Q2, whereas Q3 is again not distinguishable from Q4. In velocity, Q2 have louder notes than those of Q1, and Q3 has slightly louder notes than Q4.

Key Distribution. The valence of music is often found to be related to the major-minor tonality [7]. For simplicity, we measure the tonality from the musical key. We extract the key information using the Krumhansl-Kessler algorithm [47] in the MIDI ToolBox [48]. Figure 3 shows the key distributions on 12 major-minor pitch classes for the four emotion quadrants. They assure that the major-minor tonality is an important clue in distinguishing valence. In the high valence group (Q1, Q4), the distribution is skewed to the left (major), while, in the low valence group (Q2, Q3), the trend is the opposite. We also measured the distance between the key distributions using the Jensen-Shannon divergence, which has the property of symmetry. Table 3 summarizes the pair-wise distances, indicating that the valence difference group has a larger difference than the arousal difference group.

5. MUSIC EMOTION RECOGNITION

We report our baseline research on both symbolic- and audio-domain emotion recognition using EMOPIA, which defines the task as classifying a song into four categories. The clips in EMOPIA are divided into train-validation-test splits with the ratio of 7:2:1 in a stratified manner.

Symbolic-domain Classification. We evaluate two methods: one is based on hand-crafted features with a simple classifier, and the other is on a deep neural network model. For the former, we use the analysis features in the previous section and a logistic regression classifier as a baseline. Specifically, we use average values of note density, note

Model	4Q	A	V
Logistic regression	.581	.849	.651
LSTM-Attn [49]+MIDI-like [39]	.684	.882	.833
LSTM-Attn [49]+REMI [34]	.615	.890	.746

Table 4. Symbolic-domain classification performance.

Model	4Q	A	V
Logistic regression	.523	.919	.558
Short-chunk ResNet [50]	.677	.887	.704

Table 5. Audio-domain classification performance.

length, velocity, and represent the key as a one-hot vector. For the latter, we use two different symbolic note representation methods introduced in Section 3.4, MIDI-like [39] and REMI [34]. For the learning model, we use the combination of bidirectional LSTM and a self-attention module, or *LSTM-Attn* for short, proposed originally for sentiment classification in NLP [49]. The LSTM extracts temporal information from the MIDI note events, while the self-attention module calculates different weight vectors over the LSTM hidden states with multi-head attentions. The weighted hidden states are finally used for classification.

Audio-domain Classification. We evaluate two audio-domain classification methods in a similar manner to the symbolic-domain ones. In the first method, we use an average of 20 dimensions of mel-frequency cepstral coefficient (MFCC) vectors and a logistic regression classifier. In the second method, we use the short-chunk ResNet following [50], which is composed of 7-layer CNNs with residual connections. The output is summarized as a fixed 128-dimensional vector through max pooling, which is followed by two fully connected layers with the ReLU activation for classification. The input audio to the short-chunk ResNet is 3-second excerpts represented as a log-scaled mel-spectrogram with 128 bins with 1024-size FFT (Hanning window), and 512 size hop at 22,050 Hz sampling rate. We randomly sample three seconds of the audio chunk as an input size to the classifier.

Evaluation. We calculate 4-way classification accuracy over the four different emotion quadrants and 2-way classification accuracy over either arousal and valence. Tables 4 and 5 show the results in the symbolic and audio domains, respectively. Except for arousal, we can see that the deep learning approaches generally outperform the logistic regression classifiers using hand-crafted features. In audio domain arousal classification, MFCC vectors averaging the entire song sequence showed better performance than the deep learning approach with 3-second input. It seems that wider input sequence has the strength in emotion recognition. In both domains, valence classification is a more difficult task compared to arousal classification. For valence classification, MIDI-domain classifiers yield better result

than audio-domain classifiers (0.883 vs. 0.704). Among the two token representations, MIDI-like seems to outperform REMI for valence classification.

6. EMOTION-CONDITIONED GENERATION

We build the Transformer and LSTM models for emotion-conditioned symbolic music generation using EMPOIA. For the former, we adopt the Compound Word Transformer [35], the state-of-the-art in unconditioned symbolic music generation. We employ the CP+emotion representation presented in Section 3.4 as the data representation.

For the LSTM model, we consider the approach proposed by Ferreira *et al.* [10], which represents the state-of-the-art in emotion-conditioned music generation. Our implementation follows that described in [10], with the following differences: 1) train on EMOPIA rather than VG-MIDI; 2) use 512 neurons instead of 4,096 due to our limited computational resource; 3) use the same linear logistic regression layer for classification but we classify four classes instead of two.

As the size of EMOPIA might not be big enough, we use additionally the *AILabs1k7* dataset compiled by Hsiao *et al.* [35] to pre-train the Transformer. *AILabs1k7* contains 1,748 samples and is also pop piano music, but it does not contain emotion labels. Most of the clips in *AILabs1k7* are longer than EMOPIA, so to keep the consistency of the input sequence length, the length of the token sequence is set to be 1,024. We pre-train the Transformer with $1e-4$ learning rate on *AILabs1k7*, take the checkpoint with negative log-likelihood loss 0.30, and then fine-tune it on EMOPIA with $1e-5$ learning rate. During pre-training, the emotion token is always set to be “ignore,” while in fine-tuning it is set to the emotion of that sample.

Evaluation. We use the following three sets of metrics.

- **Surface-level objective metrics.** We use the following three metrics proposed by Dong *et al.* [43] to evaluate whether the generated samples fit the training data: pitch range (PR), number of unique pitch classes used (NPC), and number of notes being played concurrently (POLY). We use MusPy [18] to compute these metrics.
- **Emotion-related objective metrics.** Since both REMI and CP adopt a beat-based representation of music, we employ the LSTM-Attn+REMI emotion classifier (cf. Section 5) here to quantify how well the generation result is influenced by the emotion condition. We first use the generative model to generate 100 samples per class, and use the assigned label as the target class of the sample. The trained classifier is then used to make prediction on the generated samples. Similar to the classification task, apart from 4Q classification, we also conduct 2-way classification of both Arousal and Valence aspect.
- **Subjective metrics.** As the classifiers are not 100% accurate, we also resort to a user survey to evaluate the emotion-controllability of the models. Specifically, we deploy an online survey to collect responses to the music generated by different models. A subject has to lis-

Model	Objective metrics						Subjective metrics		
	PR	NPC	POLY	4Q	A	V	Humanness	Richness	Overall
EMOPIA (i.e., real data)	51.0	8.48	5.90	—	—	—	—	—	—
LSTM+GA [10]	59.1	9.27	3.39	.238	.500	.498	2.59±1.16	2.74±1.12	2.60±1.07
CP Transformer [35]	53.4	9.20	3.48	.418	.690	.583	2.61±1.03	2.81±1.03	2.78±1.03
CP Transformer w/ pre-training	49.6	8.54	4.40	.403	.643	.590	3.31±1.18	3.22±1.23	3.26±1.15

Table 6. Performance comparison of the evaluated models for emotion-conditioned symbolic music generation in *surface-level objective metrics* (Pitch Range, Number of Pitch Classes used, and POLYphony; the closer to that of the real data the better), *emotion-related objective metrics* (4Q classification, Arousal classification, Valence classification; the higher the better), and *subjective metrics* (all in 1–5; the higher the better); bold font highlights the best result per metric.

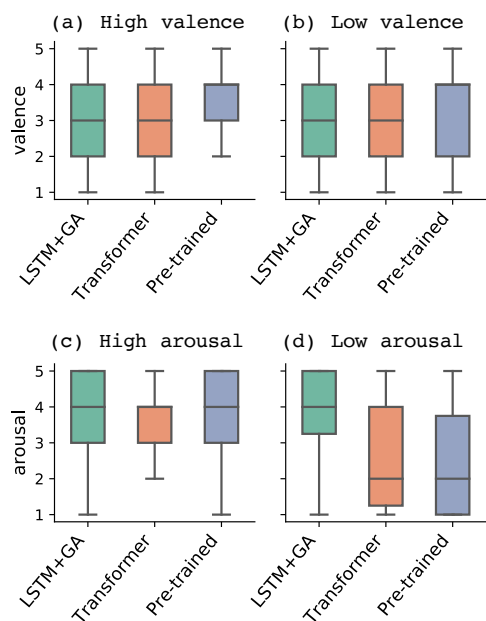


Figure 4. The subjective emotional scores (in 1–5) for the generative results when the target emotion is (a) high valence, (b) low valence, (c) high arousal, (d) low arousal.

ten to 12 random-generated samples, one for each of the three models and each of the four emotion classes, and rate them on a five-point Likert scale with respect to 1) Valence: is the audio negative or positive; 2) Arousal: is low or high in arousal; 3) Humanness: how well it sounds like a piece played by human; 4) Richness: is the content interesting; and, 5) Overall musical quality. In total 25 subjects participated in the survey.

Table 6 tabulates some of the results. We see that the CP Transformer with (‘w/’) pre-training performs the best in most of the objective metrics and the three subjective metrics listed here. Nevertheless, the scores of the CP Transformer with pre-training in the three emotion-related objective metrics are much lower than that reported in Table 4, suggesting that either the generated pieces are not emotion-laden, or the generated pieces are too dissimilar to the real pieces to the classifier.

Figure 4 shows the human assessment of the emotion-controllability of the models. To our surprise, while the CP

Transformer with pre-training does not score high in the emotion related objective metrics, the subjective test shows that it can actually control the emotion of the generated pieces to a certain extent, better than the two competing models. In particular, the valence of the samples generated by the Transformer with pre-training has a median rating of 4 when the goal is to generate positive-valence music (i.e., Figure 4(a)), while the scores of the other two models are around 3. Moreover, the arousal of the samples generated by the Transformer with pre-training has a median rating of 4 when the goal is to generate high-arousal music (Figure 4(c)), which is higher than that of the non pre-trained Transformer. This suggests that the LSTM-Attn classifier employed for computing the emotion-related objective metrics may not be reliable enough to predict the emotion perceived by human, and that the Transformer with pre-training is actually effective in controlling the emotion of the music it generates to certain extent. But, the model seems not good enough for the cases of generating low-valence (i.e., negative) music, as shown in Figure 4(b).

7. CONCLUSION

In this paper, we have proposed a new public dataset EMOPIA, a medium-scale emotion-labeled pop piano dataset. It is a multi-modal dataset that contains both the audio files and MIDI transcriptions of piano-only music, along with clip-level emotion annotations in four classes. We have also presented prototypes of models for clip-level music emotion classification and emotion-based symbolic music generation trained on this dataset, using a number of state-of-the-art models in respective tasks. The result shows that we are able to achieve high accuracy in both four-quadrant and valence-wise emotion classification, and that our Transformer-based model is capable of generating music with a given target emotion to a certain degree.

In the future, in-depth importance analysis can be conducted to figure out features that are important for emotion classification, and to seek ways to incorporate those features to the generation model. Many ideas can also be tried to further improve the performance of emotion conditioning, e.g., the Transformer-GAN approach [51].

We share not only the dataset itself but the code covering all our implemented models in a GitHub repo. We hope that researchers will find this contribution useful in future emotion-related MIR tasks.

8. ACKNOWLEDGEMENT

This research is supported by Year 2021 Culture Technology R&D Program by Ministry of Culture, Sports and Tourism and Korea Creative Content Agency (Project Name: Research Talent Training Program for Emerging Technologies in Games, Project Number: R2020040211, Contribution Rate: 50%)

9. REFERENCES

- [1] Y. Kim *et al.*, “Music emotion recognition: A state of the art review,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2010.
- [2] Y.-H. Yang and H. H. Chen, *Music Emotion Recognition*. CRC Press, 2011.
- [3] J. S. Gómez-Cañón, E. Cano, P. Herrera, and E. Gómez, “Joyful for you and tender for us: the influence of individual characteristics and language on emotion labeling and classification,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [4] J. S. Gómez-Cañón, E. Cano, Y.-H. Yang, P. Herrera, and E. Gómez, “Let’s agree to disagree: Consensus entropy active learning for personalized music emotion recognition,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2021.
- [5] 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,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [6] R. Panda, R. Malheiro, B. Rocha, A. Oliveira, and R. P. Paiva, “Multi-modal music emotion recognition: A new dataset, methodology and comparative analysis,” in *Proc. Int. Symposium on Computer Music Multidisciplinary Research*, 2013.
- [7] R. Panda, R. M. Malheiro, and R. P. Paiva, “Audio features for music emotion recognition: A survey,” *IEEE Trans. Affective Computing*, pp. 1–20, 2020.
- [8] S. Chowdhury, A. Vall, V. Haunschmid, and G. Widmer, “Towards explainable music emotion recognition: The route via mid-level features,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2019.
- [9] C. Cancino-Chacón, S. Peter, S. Chowdhury, A. Aljanaki, and G. Widmer, “On the characterization of expressive performance in classical music: First results of the con espresione game,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [10] L. N. Ferreira and J. Whitehead, “Learning to generate music with sentiment,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2019.
- [11] K. Chen, C.-I. Wang, T. Berg-Kirkpatrick, and S. Dubnov, “Music SketchNet: Controllable music generation via factorized representations of pitch and rhythm,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [12] D. Makris, K. R. Agres, and D. Herremans, “Generating lead sheets with affect: A novel conditional seq2seq framework,” in *Proc. Int. Joint Conf. Neural Networks*, 2021.
- [13] R. Madhok, S. Goel, and S. Garg, “SentiMozart: Music generation based on emotions,” in *Proc. Int. Conf. Agents and Artificial Intelligence*, 2018.
- [14] A. M. Grimaud and T. Eerola, “EmoteControl: an interactive system for real-time control of emotional expression in music,” *Personal and Ubiquitous Computing*, 2020.
- [15] T.-P. Chen, S. Fukayama, M. Goto, and L. Su, “Chord Jazzification: Learning Jazz interpretations of chord symbols,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [16] H. H. TAN and D. Herremans, “Music Fadernets: Controllable music generation based on high-level features via low-level feature modelling,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [17] Z. Wang, K. Chen, J. Jiang, Y. Zhang, M. Xu, S. Dai, X. Gu, and G. Xia, “POP909: A pop-song dataset for music arrangement generation,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [18] H.-W. Dong, K. Chen, J. McAuley, and T. Berg-Kirkpatrick, “MusPy: A toolkit for symbolic music generation,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [19] A. McLeod, J. Owers, and K. Yoshii, “The MIDI Degradation Toolkit: Symbolic music augmentation and correction,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2020.
- [20] C. Hawthorne *et al.*, “Onsets and Frames: Dual-objective piano transcription,” *arXiv preprint arXiv:1710.11153*, 2018.
- [21] Q. Kong, B. Li, X. Song, Y. Wan, and Y. Wang, “High-resolution piano transcription with pedals by regressing onsets and offsets times,” *arXiv preprint arXiv:2010.01815*, 2020.
- [22] C. Hawthorne *et al.*, “Enabling factorized piano music modeling and generation with the MAESTRO dataset,” in *Proc. Int. Conf. Learning Representations*, 2019.
- [23] K. W. Cheuk, D. Herremans, and L. Su, “ReconVAT: A semi-supervised automatic music transcription framework for low-resource real-world data,” in *Proc. ACM Multimedia*, 2021.

- [24] D. Bognadov, A. Porter, P. Tovstogan, and M. Won, “Mediaeval 2019: Emotion and theme recognition in music using Jamendo,” in *MediaEval 2019 Workshop*.
- [25] A. Alajanki, Y.-H. Yang, and M. Soleymani, “Benchmarking music emotion recognition systems,” *PLOS ONE*, 2016.
- [26] M. Defferrard, K. Benzi, P. Vandergheynst, and X. Bresson, “FMA: A dataset for music analysis,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2017.
- [27] R. Bittner, J. Salamon, M. Tierney, M. Mauch, C. Cannam, and J. Bello, “MedleyDB: A multitrack dataset for annotation-intensive mir research,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2014.
- [28] J. Fan, M. Thorogood, and P. Pasquier, “Emo-Soundscapes: A dataset for soundscape emotion recognition,” in *Proc. Int. Conf. Affective Computing and Intelligent Interaction*, 2017.
- [29] J. Fan, Y.-H. Yang, K. Dong, and P. Pasquier, “A comparative study of Western and Chinese classical music based on soundscape models,” in *Proc. Int. Conf. Acoustics, Speech, and Signal Processing*, 2020.
- [30] M. Soleymani, M. Caro, E. Schmidt, C.-Y. Sha, and Y.-H. Yang, “1000 songs for emotional analysis of music,” in *Proc. ACM Int. Workshop on Crowdsourcing for Multimedia*, 2013.
- [31] J. Fan, K. Tatar, M. Thorogood, and P. Pasquier, “Ranking-based emotion recognition for experimental music,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2017.
- [32] J. Russell, “A circumplex model of affect,” *Journal of Personality and Social Psychology*, vol. 39, pp. 1161–1178, December 1980.
- [33] A. Vaswani *et al.*, “Attention is all you need,” in *Proc. Advances in Neural Information Processing Systems*, 2017.
- [34] Y.-S. Huang and Y.-H. Yang, “Pop Music Transformer: Beat-based modeling and generation of expressive pop piano compositions,” in *Proc. ACM Multimedia*, 2020.
- [35] W.-Y. Hsiao, J.-Y. Liu, Y.-C. Yeh, and Y.-H. Yang, “Compound Word Transformer: Learning to compose full-song music over dynamic directed hypergraphs,” in *Proc. AAAI*, 2021.
- [36] J. Grekow and Z. W. Raś, “Detecting emotions in classical music from MIDI files,” in *Proc. Int. Symposium on Methodologies for Intelligent Systems*, 2009.
- [37] Y. Lin, X. Chen, and D. Yang, “Exploration of music emotion recognition based on MIDI,” in *Proc. Int. Society for Music Information Retrieval Conf.*, 2013.
- [38] K. Zhao, S. Li, J. Cai, H. Wang, and J. Wang, “An emotional symbolic music generation system based on LSTM networks,” in *Proc. IEEE Information Technology, Networking, Electronic and Automation Control Conf.*, 2019.
- [39] S. Oore, I. Simon, S. Dieleman, D. Eck, and K. Simonyan, “This time with feeling: Learning expressive musical performance,” *Neural Computing and Applications*, 2018.
- [40] C. Cannam, C. Landone, and M. Sandler, “Sonic Visualiser: An open source application for viewing, analysing, and annotating music audio files,” in *Proc. ACM Multimedia*, 2010, pp. 1467–1468, <https://www.sonicvisualiser.org/>.
- [41] L.-L. Balkwill and W. Thompson, “A cross-cultural investigation of the perception of emotion in music: Psychophysical and cultural cues,” *Music Perception*, vol. 17, pp. 43–64, October 1999.
- [42] A. Roberts, J. Engel, C. Raffel, C. Hawthorne, and D. Eck, “A hierarchical latent vector model for learning long-term structure in music,” in *Proc. Int. Conf. Machine Learning*, 2018.
- [43] H.-W. Dong, W.-Y. Hsiao, L.-C. Yang, and Y.-H. Yang, “MuseGAN: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment,” in *Proc. AAAI*, 2018.
- [44] N. S. Keskar *et al.*, “CTRL - a conditional Transformer language model for controllable generation,” *arXiv preprint arXiv:1909.05858*, 2019.
- [45] P. Juslin and E. Lindstorm, “Musical expression of emotions: Modelling listeners’ judgements of composed and performed features,” *Music Analysis*, vol. 29, pp. 334 – 364, 2011.
- [46] S. R. Livingstone, R. Muhlberger, A. R. Brown, and W. F. Thompson, “Changing musical emotion: A computational rule system for modifying score and performance,” *Computer Music Journal*, vol. 34, no. 1, pp. 41–64, 2010.
- [47] C. L. Krumhansl, *Cognitive foundations of musical pitch*. Oxford University Press, 2001.
- [48] T. Eerola and P. Toiviainen, *MIDI Toolbox: MATLAB Tools for Music Research*. Jyväskylä, Finland: University of Jyväskylä, 2004.
- [49] Z. Lin, M. Feng, C. N. d. Santos, M. Yu, B. Xiang, B. Zhou, and Y. Bengio, “A structured self-attentive sentence embedding,” in *Proc. Int. Conf. Learning Representations*, 2017.
- [50] M. Won, A. Ferraro, D. Bogdanov, and X. Serra, “Evaluation of CNN-based automatic music tagging models,” in *Proc. Sound and Music Conf.*, 2020.
- [51] A. Muhamed *et al.*, “Symbolic music generation with Transformer-GANs,” in *Proc. AAAI*, 2021.