

SIX DRAGONS FLY AGAIN: REVIVING 15TH-CENTURY KOREAN COURT MUSIC WITH TRANSFORMERS AND NOVEL ENCODING

Danbinaerin Han^{1*}
Hannah Park⁴

Mark Gotham²
Sihun Lee³

Dongmin Kim³
Dasaem Jeong⁴

¹ Graduate School of Culture Technology, KAIST, Daejeon, South Korea

² Department of Digital Humanities, King’s College London, UK

³ Dept. of Artificial Intelligence, ⁴ Dept. of Art & Technology, Sogang University, Seoul, South Korea

naerin71@kaist.ac.kr, mark.gotham@kcl.ac.uk, {dmkim, hannah, sihunlee, dasaemj}@sogang.ac.kr

ABSTRACT

We introduce a project that revives a piece of 15th-century Korean court music, *Chihwapyeong* and *Chwipunghyeong*, composed upon the poem *Songs of the Dragon Flying to Heaven*. One of the earliest examples of *Jeongganbo*, a Korean musical notation system, the remaining version only consists of a rudimentary melody. Our research team, commissioned by the National Gugak (Korean Traditional Music) Center, aimed to transform this old melody into a performable arrangement for a six-part ensemble. Using *Jeongganbo* data acquired through bespoke optical music recognition, we trained a BERT-like masked language model and an encoder-decoder transformer model. We also propose an encoding scheme that strictly follows the structure of *Jeongganbo* and denotes note durations as positions. The resulting machine-transformed version of *Chihwapyeong* and *Chwipunghyeong* were evaluated by experts and performed by the Court Music Orchestra of National Gugak Center. Our work demonstrates that generative models can successfully be applied to traditional music with limited training data if combined with careful design.

1. INTRODUCTION

Six dragons fly on the east land; every endeavour is a heavenly blessing. This is the first line of lyrics in *Yongbieocheonga*, the first text written in the Korean alphabet (Hangul, 한글). Sejong the Great, one of the most respected figures in Korean history, invented and introduced Hangul in 1446. In addition to this remarkable achievement, he ordered scholar-officials to write *Yongbieocheonga*, and composed music to accompany the lyrics. Three other pieces composed at the time are *Yeo-Min-Lak*, *Chi-Hwa-Pyeong* and *Chwi-Pung-Hyeong*.

* Work mainly done during her master’s at Sogang University

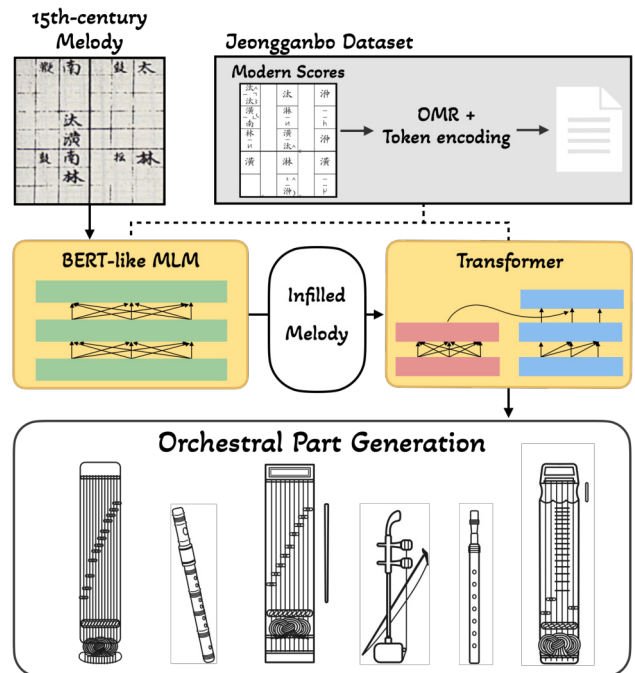


Figure 1: Overview of the proposed research framework

These compositions are still preserved in the *Veritable Records of Sejong*, which is the oldest surviving musical score in Korea [1]. More detailed information is available here [2].

Among these three pieces, only *Yeominlak* is handed down to the present day, while the other two are no longer performed. The National Gugak Center¹, which is the primary organization dedicated to the preservation and development of traditional music, commissioned the task of reconstructing these two pieces in a performable format using artificial intelligence systems. Given a simple melody of 512 *gaks* (measures) of *Chihwapyeong* or 132 *gaks* of *Chwipunghyeong*, the system must generate scores for six different instruments.

Our solution encompasses a wide range of tasks in the field of music information retrieval—constructing a specialized dataset, optical music recognition, designing a

¹ *Gugak* (국악) is the Korean term for traditional music

domain-specific encoding scheme, training models with limited data, and generating music of concert-level quality. In this paper, we present in detail the different frameworks used in the project: two types of transformer-based models; a symbolic dataset of Korean court music acquired through optical music recognition; and a novel “Jeonggan-like” encoding method that notates monophonic melody by combining notes’ position and pitch, along with a beat counter that informs the transformer the temporal position. The effectiveness of the proposed techniques was validated through quantitative metrics and subjective evaluation by experts from the National Gugak Center. Finally, we introduce a web demo that allows users to examine and generate traditional Korean court music interactively.

This project has significance not only for cultural preservation but also for wider considerations in machine learning and music generation. One of the many benefits to be had from the inter-cultural study of music is the different perspectives expressed in ‘the music itself’ as well as any notational and/or theoretical traditions that go alongside it. As presented in previous research [3], the encoding of music makes significant differences in machine learning tasks. In thinking through different ways of digitally encoding music, we stand to learn a great deal from the various syntaxes that have been used in diverse traditional contexts.

2. RELATED WORKS

Recent advances in neural network-based music generation have resulted in much artistic output. Since 2020, the *AI Music Generation Challenge* [4] has been held annually, focusing on generating songs in the style of Irish and Swedish folk music. This event has allowed for exploration of new methods for generation and evaluation of traditional music through the means of deep learning models.

The *Beethoven X project* [5] utilized neural networks to learn Beethoven’s compositional style and complete his unfinished 10th Symphony. The resulting work has been performed by an orchestra—a project outline similar to that of ours.

Attempts at automatic generation have been made for traditional music from beyond the West, including Persia [6] and China [7]. The limited progress in such areas is often due to the distinctive traditional musical systems that demand deep understanding and unique methodologies. Such idiosyncrasies put much interest and meaning in the computational research of traditional music, since it can present new methods and perspectives to the field as a whole, while also helping preserve diverse musical heritages.

3. JEONGGANBO DATASET

3.1 Jeongganbo Notation

As depicted in Figure 1, Korean court music is performed on a variety of instruments, including plucked string instruments (*Gayageum* and *Geomungo*), bowed string instruments (*Haegum* and *Ajaeng*), and wind instruments




Figure 2: An example of Jeongganbo in the original notion (below) and a broadly equivalent conversion to Western classical notion (above). Dashed lines are part of neither notation and added simply to clarify the temporal alignment between the two systems.




Figure 3: Jeonggan-like encoding position labels

(*Daegum* and *Piri*), among others. These instruments are played together in a heterophonic texture, with each instrument employing its distinctive playing techniques and ornamentations.

Much of Korean court music is written in *Jeongganbo*, a traditional musical notation system. *Jeongganbo* is recognized as the first system in East Asia capable of simultaneously representing both pitch and duration of notes [8,9]. This versatility has been instrumental in passing down court music throughout history [10].

Jeongganbo uses grid-divided boxes (*Jeonggans*) as the basic unit of time. The *number* of characters (notes) and their *position* within each *jeonggan* varies to denote rhythm. Figure 2 provides an example passage, and figure 3 provides a schematic overview of possible positions.

Here, we provide a broad introduction to this rhythmic notation system in quasi-Western musical theoretic language. Each *jeonggan* is broadly equivalent to a beat. If a *jeonggan* features only one character, this note event starts at the beginning of the beat and lasts the beat’s full duration. The first box (‘0’) in figure 3 is in this form as is the second *jeonggan* of figure 2 where the ‘compound beats’ correspond to the duration ♩. (in this case for the note B♭4). At the next metrical level we have the ‘column’ division of the ‘rows’. This number of ‘rows’ relates broadly to the top level division of the beat. The use of three vertically stacked characters refers to 3 equal divisions of this beat (here, 3 x ♩s). For example, in figure 3, the numbers 4–9 feature a 3-part division of the ♩ beat into 3 x ♩s (positions 4, 6, 8), and a 2x division of those ♩s (e.g., 4–5). If the following *jeonggan* is empty, the previously played note is sustained.

Playing techniques and ornamentations called *sigimsae* are sometimes notated for each instrument. When *sigimsae* are placed to the right of notes, they serve as ornamentations or embellishments for the corresponding note; when written on their own, they indicate timed instructions to play a specific note or musical phrase. For convenience,

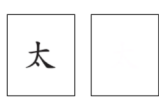

Jeonggan		
JG-like	:0 太	:1 太 ̂ :7 □ :3 仲
REMI-like	0 太	0 太 ̂ 1/2 □ 2/3 仲
ABC-like	太 6	太 ^ 1.5 ▮ 0.5 仲 1

Figure 4: Comparison between encoding schemes

the example score is notated horizontally, but in practice, the score page is read from top to bottom and right to left. A line in *Jeongganbo* can consist of anything from four to twenty beats, with each line representing a phrase unit.

3.2 Machine Readable Dataset

We have constructed a dataset of 85 pieces by applying optical musical recognition (OMR) to all compositions available within the manuscripts published by the National Gugak Center. The manuscripts cover the entire repertoire of remaining Korean court music². OMR was necessary since the scores are only provided as PDF images and the semantic data is unavailable. We implemented and trained an encoder-decoder transformer with CNN by synthesizing various *Jeonggan* images in a rule-based approach [11]. In total, the dataset comprises 28 010 jeonggans across 85 pieces. When counting each instrument part independently, the combined total amounts to 141 820 jeonggans. Out of 90 pieces notated in jeongganbo for ensembles of at least two different instruments in the published manuscripts, we excluded 5 pieces that have discrepancies in the total number of jeonggans across instruments.

4. JEONGGAN-LIKE ENCODING

In the field of symbolic music generation for Western monophonic and polyphonic music, encoding schemes such as ABC notation, which denotes pitch and duration separately, are effective and prevalent [12, 13]. However, when it comes to Korean court music, whose heterophonic structure is a defining characteristic, it is crucial that the intricate alignment of different melodies be well-represented in encoding. The genre also exhibits prolonged notes and considerable variations in note lengths, which proves to be a challenge for learning algorithms, especially when data is limited.

These distinct musical qualities call for a specialized encoding scheme; for this, we propose *Jeonggan (JG)-like encoding*, which closely follows the positional notation of *Jeongganbo*. This symbolic music encoding method is modeled to inherently reflect the composition and notation style of traditional Korean court music.

The detailed rules of encoding are as follows. The boundary of a *Jeonggan* is designated as a bar (|) to-

² The term “court music” used in this paper originally refers to *Jeong-ak*. Jeong-ak includes not only court music but also salon music and military music. However, for readability, we use “court music” here.

ken. Change of measure (called *Gak*) is indicated by a line break ($\backslash n$). As illustrated in Figure 3, the position of each note is denoted by a number between 0 and 15, after which the pitch symbol follows.

Ornamentations (*sigimsae*) can either have a duration or not. *Sigimsae* with duration, such as the ‘ \neg ’ symbol in Figure 4, are handled in the same way as pitch symbols. *Sigimsae* without duration such as ‘ \wedge ’, which appear at the side of the pitch character, are placed after the corresponding pitch symbol.

There are several advantages that we can expect to gain from using *JG-like* encoding. First, with position-based encoding, the duration-related vocabulary is limited to just 16 entries. In contrast, duration-based encoding schemes require learning each duration token as a separate entry, resulting in a significantly larger vocabulary. Additionally, rather than determining the length of a note with a single calculation, JG-like encoding allows for the flexible adjustment of note lengths during inference via combination of *jeonggan* boundary and position tokens. This enables generation of music that is more adaptable to the time step and takes into account the sequence of the input source, which can be expected to result in more dynamic and context-aware music generation.

4.1 Other Possible Encodings

REMI (revamped MIDI-derived events) [14] first proposed the usage of beat-position feature rather than time-shifting to encode temporal position. We also experiment with REMI-like encoding which adopts three token types: beat position, new beat (instead of new measure), and pitch tokens. We intentionally design REMI-like and JG-like encoding to share the same structure and result in the same number of tokens for a given melody. They differ in that JG encoding provides intra-JG position, while REMI encoding provides the beat position of the note. According to the position labels shown in Figure 3, any of [0, 1, 4, 10, 12] can correspond to beat position 0. However, in JG-like encoding, each occurrence of position tokens limits the possibilities of subsequent ones. For instance, a position token of 0 implies that no more notes will occur in the same *jeonggan*, and if the first note is 1, one or more additional notes should follow with values of 2-3 or 6-9. In contrast, in REMI-like encoding, any offset value can follow a beat position of 0. To examine the impact of this position-based logic on the generation process, we use REMI-like encoding as our first baseline for comparison.

As a second baseline, we implement an ABC-like encoding scheme that does not have a separate bar token and encodes each note as a combination of pitch and duration values. Note that we do not omit duration tokens that are equal to unit length as ABC encoding typically does.

5. ORCHESTRAL PART GENERATION

5.1 Transformer Sequence-to-sequence Model

We implement an encoder-decoder transformer [15] model to generate melodies for different instruments based on a




Figure 5: Orchestral part generation

given instrument’s melody, leveraging its ability to learn long-term dependencies. Unlike RNN-based models, the transformer calculates relationships between all elements in a sequence via the self-attention mechanism, enhancing its capability in symbolic music generation [16, 17, 18]. The model consists of an encoder that processes the input sequence and a decoder for generating the output sequence. Our objective is to generate melodies that synchronize with the input melody across musically equivalent phrases; self- and cross-attention within the model enable understanding of musical context at measure and bar levels, capturing the repeating structure of melodies and accents prominent in traditional Korean court music.

5.2 Beat Counter

Instead of sinusoidal [15] or learned [19] positional embedding commonly utilized in transformer-based models, we implemented a ‘beat counter’ embedding that provides information about temporal position.

For a model to learn to ‘parse’ semantic position only from the tokens’ sequential position is challenging, if not impossible, with limited training data and a small number of transformer layers. Therefore, we explicitly encode the musical position of each symbol as a combination of measure index, beat index, and sub-beat index (in-*jeonggan* position) as shown in Figure 5. This information is summed into note embedding, just like positional encoding of transformer [15].

As previous research of PopMAG [20] demonstrated, metrical position embeddings can replace the positional encoding of transformers in symbolic music. A minor difference between PopMAG and our approach is that the model predicts only the appearance of new measures or new beats without the index of them, and that the new beat can be used for elongating the duration of previous note.

The same idea of embedding the beat counting has been

previously applied to RNN-based Irish melody generation in ABC format [21], while its advantage was not properly evaluated. A similar idea, using metrical position instead of or along with absolute token position, has also been applied to transformer architecture [20, 22, 23]. However, our results presented in Section 6.3 demonstrate that this beat counter embedding is essential for making the model properly understand the musical contents.

6. EXPERIMENT AND RESULTS

6.1 Training

We split the *Jeongganbo* dataset into three subsets: 75 pieces for training, 5 for validation, and 5 for testing. Each piece contains melodies for up to 6 instruments. The sequence-to-sequence model takes 4 measures of melody, each from a randomly selected number of instruments, as input to the encoder; and given a target instrument condition, it generates the corresponding 4 measures of melody for the target instrument. Note that the number of beats in a single measure is at least 4 or to a maximum of 20 in our dataset.

The transformer encoder and decoder both consist of 6 layers, 4 attention heads, and a hidden dimension size of 128 with dropout of 0.2. We train for 35 000 updates across 300 epochs using negative log-likelihood loss. We also employ mixed precision training [24] to enhance performance and efficiency. We utilize the Adam optimizer with an initial learning rate of 0.001 and apply a cosine learning rate scheduler with 1000 warmup steps. Using a batch size of 16, training can be conducted on a single Nvidia RTX A6000 GPU.

6.2 Evaluation Metrics

6.2.1 Length Match Rate

As an evaluation metric, we check whether the input and output melodies share the same number of measures, a consistency necessitated by our task. Since the length of a measure can change in the middle of a piece, this metric serves as an indicator of the model’s ability to capture the musical context of the input melody and accordingly generate a musically complete melody. We measure *length match rate* as the percentage of generated melodies whose number of *jeonggans*, after decoding the output tokens, matches that of the input melody.

6.2.2 F₁-Score

Regarding the generation task as one with a fixed answer, we can measure the accuracy of the generated melody by directly comparing it with the ground-truth target melody. Thus, as a general accuracy metric, we calculate the F₁-score of predicted notes, where only the notes with the exact same onset position and pitch are counted as correct. To make a fair comparison between encoding methods, note onset positions in JG-like encoding were converted to those in the REMI-like format. Ornamentations without duration were not counted.

	<i>Piri to Geom.</i>		<i>Every to Daeg.</i>	
	len-mat	F1	len-mat	F1
JG-like	0.942	0.679	1.0	0.614
REMI-like	0.923	0.567	1.0	0.532
ABC-like	1.0	0.704	0.903	0.542
JG w/o Counter	0.135	0.043	0.269	0.052
REMI w/o Counter	0.269	0.081	0.192	0.039
ABC w/o Counter	0.403	0.090	0.115	0.016

Table 1: Quantitative evaluation results

6.3 Results and Discussion

In our sequential generation process, melody for the instrument geomungo, characterized by its low pitch range and simple melodies, is the first to be generated from the initial piri melody. The daegeum, typically featuring the most complex and nuanced melodies among the six instruments, is the last in line. Table 1 displays the results of objective evaluation, specifically focusing on geomungo and daegeum.

For generation of geomungo melodies, ABC-like encoding yields the best results. This appears to be due to the simple and regular melodic structure of the geomungo which fits in well with ABC-like encoding. On the other hand, in the task of generating daegeum melodies, JG-like encoding achieves higher F₁-scores. This indicates that JG-like encoding outperforms other methods in generating complex and varied melodies. We also discover that as rhythmic complexity increases, the measure length match rate of ABC-like encoding decreases.

To examine the effectiveness of the beat counter technique, we compare our model that incorporates beat counter with a baseline model that instead employs absolute position embedding [19], a technique commonly used in symbolic music generation.

The results in the lower part of Table 1 show that the models without beat counter fail to generate melodies with appropriate lengths. The problem is less severe in ABC-like encoding, as processing accumulating duration tokens can be easier than counting *jeonggan* boundaries. This demonstrates the efficacy of the beat counter technique in JG-like encoding, and its ability to replace traditional positional encoding.

7. 15TH CENTURY MELODY TRANSFORMATION

To generate an entire ensemble score using our method, we require an initial input melody with a specified instrument. However, the remaining 15th-century score of *Chihwapyeong* and *Chwipunghyeong* only provide a single melody without any mention of instruments. It also features rhythmic groupings of eight beats, which is rare in court music that is played today. We therefore need to transform the old melody for a specific instrument used in court music; to maintain the outline of the original melody while achieving plausible transformation, we train a masked language model on our *Jeongganbo* dataset be-

fore infilling the 15th-century melody.

7.1 BERT-like Masked Language Model

Bidirectional Encoder Representations from Transformers (BERT) [25] is a self-supervised language representation learning model that uses a bidirectional transformer instead of a causal transformer decoder. It is trained with a masked language model (MLM) objective, where tokens in the input sentence are randomly masked and the model predicts the original vocabulary ID of said masked tokens. Because of its advantage in exploiting bidirectional context, BERT-like models have also been adapted for music audio generation [26] and symbolic music generation [27, 28] along with representation-learning purpose adaptation on symbolic music [22, 29].

7.1.1 Piano-roll-like Encoding

One of the main limitations of using a BERT-like model for generative tasks is that the sequence of given (unmasked) tokens and masked tokens has to be pre-defined. This means that one has to decide the number and position of new tokens to be inserted for a given original sequence. To avoid this, we use piano-roll-like encoding for the MLM, a technique widely employed in works on music generation with limited rhythmic patterns such as in Bach Chorales [30, 31, 32, 33]. Here, each *jeonggan* is represented as six frames, with each frame including features for symbol (pitch or *sigimsae* with duration) and for ornamentation. We also apply the aforementioned beat counter in piano-roll encoding.

7.1.2 Training with Masking

Following examples in MusicBERT [29], we train the model with masked language model objective with various masking methods: i) masking 5% of frames, ii) replacing 5% of frames, iii) masking 20% of note onsets, iv) replacing 10% of note onsets, v) erasing 10% of note onsets, vi) masking the entire 6 frames of 15% of *jeonggans*, and vii) masking 50% of ornamentations.

Though the model can be trained to handle an arbitrary number of input instruments, we only train the model with a single instrument as with our orchestration transformer, since the main intended usage of the model is to create variations of a single melody. We train a 12-layer model with the same dataset and hyperparameter settings as with the orchestration model.

7.2 Inference Procedure

For converting and performing monophonic melodies, we opt for a 30x ♩ span which equals to 10 *jeonggans*. This also corresponds to the rhythmic pattern of the 4–7th movement of *Yeominlak*. The original *Chihwapyeong* and *Chwipunghyeong* melody, which can be interpreted in an 8/8 time signature, were modified by strategically inserting empty *jeonggans* to the 5th and 7th positions, to imitate *Yeominlak*'s rhythmic pattern. Utilizing the masked language model, the modified melodies were seamlessly

transformed into a piri melody. Piri, a double-reed instrument known for its loud volume, was chosen as the main instrument for conveying the original melody due to its prominent role in contemporary court music.

As the models were all trained on 4-measure chunks, we generate the full sequence of 512 or 132 measures using a moving window, providing two measures of previously generated output as teacher-forcing inputs and generating one more measure for each four-measure input. These were applied in a similar manner to both melody transformation and orchestral part generation. Once the melody is transformed into a piri melody, we feed it to the orchestral transformer to generate parts for five other instruments. We sequentially generate for each instrument with the previously generated part as input. The final generation order is as follows: piri, geomungo, gayageum, ajaeng, haegeum, and daegeum.

Following the initial generation of melodies for all six instruments, we perform a refinement step. Here, each instrument’s melody is regenerated with the melodies of the other five as input. This additional process helps to reinforce the melodies that initially had to be generated without the context of the other instruments.

7.3 Expert Reviews

The Court Music Orchestra of the National Gugak Center performed the generated *Chihwapyeong* and *Chwipunghyeong* on the birth anniversary of King Sejong at Gyeongbokgung Palace on May 14th, 2024. They performed it again at the National Gugak Center on June 2nd, 2024 with an introduction to technical background by the authors. Due to time constraints, only partial excerpts from the entire score were performed.

The musicians gave positive opinions such as “*genre-specific rhythm and melodic flow were well-represented*” and “*the generated pieces presented ornamentation techniques and melodic progressions specialized for each instrument.*” Still, there were a few instances where notes that did not fit the scale appeared, and when notes outside the appropriate range were present, the performers had to alter or omit them or change their octave to perform the piece. However, the generated results were acknowledged to closely resemble the target style of Yeominlak. Thus, the Court Music Orchestra decided to play the pieces in a similar ensemble size to Yeominlak without further modification.

We additionally evaluate the generated scores, focusing on the effects of the refinement step. The evaluation criteria were carefully selected to assess aspects that require a deep understanding of the genre. These criteria include 1) the appropriateness of the scale and range for each instrument (*scale*), 2) the proper use of unique characteristics and ornamentations specific to each instrument (*sigimsae*), 3) the suitability of the rhythmic structure of strong and weak beats (*rhythm*), and 4) the harmony and coherence among the instruments when performed together as an ensemble (*harmony*).

Seven employees from the National Gugak Center who

	<i>No Refinement</i>	<i>With Refinement</i>
Scale	4.0 (± 0.53)	4.0 (± 0.53)
Sigimsae	3.4 (± 0.73)	4.0 (± 0.53)
Rhythm	2.9 (± 0.35)	3.3 (± 0.45)
Harmony	2.9 (± 0.83)	3.3 (± 0.70)

Table 2: the average and std of opinion scores from 7 judges for systems with and without refinement.

majored in Korean traditional music instruments or theory participated in a subjective survey. We name the pre-refinement generation results piece A, and the final output after the refinement step piece B. The evaluators were not informed of this distinction. The participants assessed the pieces for the four criteria on a 5-point scale (1-5) and provided qualitative feedback on the two compositions. The results are summarized in Table 2. These results demonstrate that the proposed refinement process effectively enhances the overall quality of the generated music, especially for *sigimsae* of each instrument.

8. CONCLUSION

Throughout this work, we explored how music generation models can resurrect ancient melodies into new compositions that meet style of current-day Korean court music.

Venturing into relatively uncharted territory, we approached each step meticulously—from data curation and parsing to model architecture design—while carefully considering the unique nuances of the musical tradition. To enhance the quality of the generated outputs, we proposed a novel encoding framework and validated its effectiveness through objective and subjective measures. This endeavour to tackle an underrepresented non-Western music genre through diverse MIR lenses hopefully expands the horizons of the field.

The *Jeongganbo* dataset and its conversion to Western staff notation in MusicXML is available online, along with other code of this project, and video recording of the performance.³ To the best of our knowledge, this will be the first dataset of machine-readable *Jeongganbo*. We believe that this dataset can significantly contribute to computational ethnomusicology beyond its usage as a training dataset for music generation demonstrated in this paper.

We also provide an interactive web demo⁴ that showcases our proposed generative model. While this project focused on reviving melodies from the 15th-century, the web demo allows users to input their own melodies and create orchestrations of Korean court music. The interactive platform enables users to directly engage with the generative model in the web browser.

We hope that this project contributes to moving closer to leveraging machine learning to make traditional music more accessible and enjoyable for modern audiences.

³ <https://github.com/MALerLab/SejongMusic>

⁴ <https://six-dragons-fly-again.site/>

9. ACKNOWLEDGEMENTS

We sincerely appreciate the National Gugak Center and its staff who supported this project, including director-general Kim Youngwoon (김영운), head of the research bureau Kim Myung-suk (김명석), research officers Park Jeonggyeong (박정경) and Han Jungwon (한정원). We are deeply grateful to the musicians of the Court Music Orchestra for their invaluable contributions and efforts to vitalize our humble results. This research was also supported by the National R&D Program through the National Research Foundation of Korea (NRF) funded by the Korean Government (MSIT) (RS-2023-00252944, Korean Traditional Gagok Generation Using Deep Learning).

10. REFERENCES

- [1] Y. Kim, "Chapter III. critical assessment : The rhythmic interpretation of jeongganbo," in *Korean Musicology Series, vol. 4*, 2010.
- [2] National Gugak Center, "Publications," <https://www.gugak.go.kr/site/program/board/basicboard/view?menuid=001003002005&pagesize=10&boardtypeid=24&boardid=13154&lang=en>, 2024.
- [3] G. Micchi, M. Gotham, and M. Giraud, "Not all roads lead to Rome: Pitch representation and model architecture for automatic harmonic analysis," *Transactions of the Int. Society for Music Information Retrieval (TISMIR)*, vol. 3, no. 1, pp. 42–54, 2020.
- [4] B. Sturm, "The Ai music generation challenge 2022: Summary and results," in *Proc. of the 4th Conference on AI Music Creativity (AIMC 2023)*, Brighton, UK, 2023.
- [5] M. Gotham, K. Song, N. Böhlefeld, and A. Elgammal, "Beethoven X: Es könnte sein!(it could be!)," in *Proc. of the 3rd Conference on AI Music Creativity (AIMC 2022)*, Online, 2022, pp. 13–15.
- [6] M. Ebrahimi, B. Majidi, and M. Eshghi, "Procedural composition of traditional Persian music using deep neural networks," in *Proc. of 5th Conference on Knowledge Based Engineering and Innovation (KBEI)*. Tehran, Iran: IEEE, 2019, pp. 521–525.
- [7] J. Luo, X. Yang, S. Ji, and J. Li, "MG-VAE: Deep chinese folk songs generation with specific regional styles," in *Proc. of the 8th Conference on Sound and Music Technology (CSMT) Revised Selected Papers*. Shanxi, China: Springer, 2020, pp. 93–106.
- [8] A. E. Gnanadesikan, *The writing revolution: Cuneiform to the internet*. John Wiley & Sons, 2008, vol. 8.
- [9] Y. Kim, "Chapter II. Korean notational systems," in *Korean Musicology Series, vol. 4*, 2010.
- [10] R. Koehler *et al.*, *Traditional music: sounds in harmony with nature*. Seoul Selection, 2015.
- [11] D. Kim, D. Han, D. Jeong, and J. J. Valero-Mas, "On the automatic recognition of jeongganbo music notation: dataset and approach," *Preprint on Research Square*, 2024.
- [12] B. Sturm, J. F. Santos, and I. Korshunova, "Folk music style modelling by recurrent neural networks with long short term memory units," in *late-breaking demo session of 16th Int. Society for Music Information Retrieval Conf.*, Málaga, Spain, 2015.
- [13] B. L. Sturm, J. F. Santos, O. Ben-Tal, and I. Korshunova, "Music transcription modelling and composition using deep learning," *arXiv preprint arXiv:1604.08723*, 2016.
- [14] Y.-S. Huang and Y.-H. Yang, "Pop music transformer: Beat-based modeling and generation of expressive pop piano compositions," in *Proc. of the 28th ACM International conference on multimedia*, New York, NY, USA, 2020, pp. 1180–1188.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. of 31st Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, USA, 2017.
- [16] C.-Z. A. Huang, A. Vaswani, J. Uszkoreit, I. Simon, C. Hawthorne, N. Shazeer, A. M. Dai, M. D. Hoffman, M. Dinculescu, and D. Eck, "Music transformer," in *Proc. of The 7th International Conference on Learning Representations*, New Orleans, LA, USA, 2019.
- [17] B. Yu, P. Lu, R. Wang, W. Hu, X. Tan, W. Ye, S. Zhang, T. Qin, and T.-Y. Liu, "Museformer: Transformer with fine- and coarse-grained attention for music generation," in *Proc. of The 36th Annual Conference on Neural Information Processing Systems*, New Orleans, LA, USA, 2022, pp. 1376–1388.
- [18] Y.-J. Shih, S.-L. Wu, F. Zalkow, M. Müller, and Y.-H. Yang, "Theme transformer: Symbolic music generation with theme-conditioned transformer," *IEEE Transactions on Multimedia*, vol. 25, pp. 3495–3508, 2023.
- [19] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin, "Convolutional sequence to sequence learning," in *Proc. of the 34th International Conference on Machine Learning*, Sydney, Australia, 2017, pp. 1243–1252.
- [20] Y. Ren, J. He, X. Tan, T. Qin, Z. Zhao, and T.-Y. Liu, "Popmag: Pop music accompaniment generation," in *Proc. of the 28th ACM Int. Conf. on Multimedia*, 2020, pp. 1198–1206.
- [21] D. Jeong, "Virtuosotune: Hierarchical melody language model," *IEIE Transactions on Smart Processing & Computing*, vol. 12, no. 4, pp. 329–333, 2023.

- [22] Z. Wang and G. Xia, “MuseBERT: Pre-training music representation for music understanding and controllable generation,” in *Proc. of the 22nd Int. Society for Music Information Retrieval Conf.*, Online, 2021, pp. 722–729.
- [23] Z. Guo, J. Kang, and D. Herremans, “A domain-knowledge-inspired music embedding space and a novel attention mechanism for symbolic music modeling,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 4, 2023, pp. 5070–5077.
- [24] P. Micikevicius, S. Narang, J. Alben, G. Diamos, E. Elsen, D. Garcia, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh *et al.*, “Mixed precision training,” *arXiv preprint arXiv:1710.03740*, 2017.
- [25] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Minneapolis, Minnesota, USA, 2019, pp. 4171–4186.
- [26] H. Flores García, P. Seetharaman, R. Kumar, and B. Pardo, “Vampnet: Music generation via masked acoustic token modeling,” in *Proc. of the 24th Int. Society for Music Information Retrieval Conf.*, Milan, Italy, 2023.
- [27] R. Dahale, V. Talwadker, P. Rao, and P. Verma, “Generating coherent drum accompaniment with fills and improvisations,” in *Proc. of the 23rd Int. Society for Music Information Retrieval Conf.*, Bengaluru, India, 2022.
- [28] L. Casini, N. Jonason, and B. L. T. Sturm, “Investigating the viability of masked language modeling for symbolic music generation in abc-notation,” in *Proc. of 13th International Conference on Computational Intelligence in Music, Sound, Art and Design*, Aberystwyth, UK, 2024, pp. 84–96.
- [29] M. Zeng, X. Tan, R. Wang, Z. Ju, T. Qin, and T.-Y. Liu, “MusicBERT: Symbolic music understanding with large-scale pre-training,” in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, Online, 2021, pp. 791–800.
- [30] F. Liang, M. Gotham, M. Johnson, and J. Shotton, “Automatic stylistic composition of bach chorales with deep LSTM,” in *Proc. of 18th Int. Society for Music Information Retrieval Conf.*, Suzhou, China, 2017, pp. 449–456.
- [31] G. Hadjeres, F. Pachet, and F. Nielsen, “DeepBach: a steerable model for Bach chorales generation,” in *Proc. of the 34th International Conference on Machine Learning*, Sydney, Australia, 2017, pp. 1362–1371.
- [32] C.-Z. A. Huang, T. Cooijmans, A. Roberts, A. Courville, and D. Eck, “Counterpoint by convolution,” in *Proc. of 18th Int. Society for Music Information Retrieval Conf.*, Suzhou, China, 2017.
- [33] E. Choi, H. Kim, J. Nam, and D. Jeong, “Teaching chorale generation model to avoid parallel motions,” in *Proc. of The 16th International Symposium on Computer Music Multidisciplinary Research (CMMR 2023)*, Tokyo, Japan, 2023.