

# MELODYT5: A UNIFIED SCORE-TO-SCORE TRANSFORMER FOR SYMBOLIC MUSIC PROCESSING

Shangda Wu<sup>1,‡</sup>      Yashan Wang<sup>1,‡</sup>      Xiaobing Li<sup>1</sup>      Feng Yu<sup>1</sup>      Maosong Sun<sup>1,2,‡</sup>

<sup>1</sup> Central Conservatory of Music, China      <sup>2</sup> Tsinghua University, China

{shangda, alexis\_wang}@mail.ccom.edu.cn,      sms@tsinghua.edu.cn

<https://github.com/sanderwood/melodyt5>

## ABSTRACT

In the domain of symbolic music research, the progress of developing scalable systems has been notably hindered by the scarcity of available training data and the demand for models tailored to specific tasks. To address these issues, we propose MelodyT5, a novel unified framework that leverages an encoder-decoder architecture tailored for symbolic music processing in ABC notation. This framework challenges the conventional task-specific approach, considering various symbolic music tasks as score-to-score transformations. Consequently, it integrates seven melody-centric tasks, from generation to harmonization and segmentation, within a single model. Pre-trained on MelodyHub, a newly curated collection featuring over 261K unique melodies encoded in ABC notation and encompassing more than one million task instances, MelodyT5 demonstrates superior performance in symbolic music processing via multi-task transfer learning. Our findings highlight the efficacy of multi-task transfer learning in symbolic music processing, particularly for data-scarce tasks, challenging the prevailing task-specific paradigms and offering a comprehensive dataset and framework for future explorations in this domain.

## 1. INTRODUCTION

In the field of artificial intelligence, symbolic music processing—including the analysis and generation of musical scores—presents a unique challenge that merges musical creativity with computational complexity. Symbolic music, which represents musical information with discrete symbols rather than continuous audio signals, facilitates the precise manipulation and analysis of elements such as melody, harmony, and rhythm. Historically, the application of AI in this area has sought not only to mimic the creative process of human composers [1–4] but also

to uncover the underlying patterns of musical composition [5–7].


Despite significant progress, the field still faces persistent limitations. One notable challenge is the prevalence of task-specific models [8–11]. These models offer benefits for specific applications but lack adaptability to the broader spectrum of symbolic music processing. This fragmentation is further compounded by the scarcity of annotated datasets [12–14], which serve as the lifeblood of deep learning models. Unlike other domains where data may be abundant and easy-to-collect, annotated symbolic music datasets are both rare and costly to produce. Without access to ample and diverse data, models struggle to generalize and may exhibit biases or limitations [15] in their analysis and generation of symbolic music.

In addressing the challenges inherent to symbolic music processing, insights from the Natural Language Processing (NLP) domain offer a promising avenue for advancement. Techniques such as transfer learning [16–18] and multi-task learning [19–21] have played a pivotal role in advancing NLP by promoting the transfer of knowledge from pre-trained language models and exploiting common patterns across various tasks. Prominent models like GPT [22], BERT [23], and T5 [24] demonstrate the efficacy of these strategies in understanding and generating language across diverse contexts. Notably, the T5 model, with its text-to-text framework, mirrors the conceptual shift necessary for symbolic music by treating all tasks as variations of converting input scores to output scores. By embracing such methodologies, which regard tasks as facets of a unified problem, we seek to develop models for symbolic music that not only excel in specific tasks but are also adaptable and proficient across a wide range of tasks.

In this paper, we introduce MelodyT5, which leverages an encoder-decoder framework to perform multiple symbolic music tasks as unified score-to-score transformations. Pre-trained on the MelodyHub dataset, which contains over 1 million task instances across seven melody-centric tasks in ABC notation, MelodyT5 overcomes the limitations of task-specific models and sparse data availability in symbolic music processing. By implementing bar patching [7, 25], MelodyT5 can handle longer sequences effectively, expanding its applicability to a wider range of tasks while maintaining computational efficiency. Our results underscore the promise of employing multi-task learning approaches in symbolic music processing, demonstrating

<sup>‡</sup> These authors contributed equally.

<sup>‡</sup> Corresponding author.



**Figure 1.** The MelodyT5 framework employs a Transformer encoder-decoder architecture with bar patching for music processing. It uses linear projection of input bar patches, fed into a patch-level Transformer encoder. The encoder output provides context for a patch-level Transformer decoder to autoregressively produce target bar features. A character-level Transformer decoder then uses these features to generate detailed characters for each bar, forming the target musical score.

the superior performance of MelodyT5 across a spectrum of tasks and providing a rich dataset for future research. The key contributions of our paper are as follows:

- MelodyT5, employing an encoder-decoder approach, redefines symbolic music processing by including multiple tasks as unified score-to-score transformations, demonstrating versatility and breaking traditional task-specific constraints.
- MelodyHub, a dataset comprising 261,900 unique melodies in ABC notation across over 1 million task instances for seven different tasks, serves as the cornerstone for effective pre-training of MelodyT5.
- Our experimental results demonstrate the efficacy of multi-task transfer learning in symbolic music, with models trained across multiple tasks outperforming those trained in isolation.

## 2. METHODOLOGY

In this section, we delve into the methodology behind MelodyT5. We first introduce the ABC notation and bar patching for music representation, then present the architectural design of MelodyT5, and finally outline the pre-training objective of our model, which focuses on score-to-score transformations as the basis for multi-task learning.

### 2.1 Data Representation

We utilize ABC notation, a concise symbolic music format, for encoding musical scores with ASCII characters. This text-based format elegantly represents musical elements like notes, rhythms, and articulations in a human-readable manner, thereby facilitating thorough music documentation. Additionally, it promotes the utilization of NLP techniques for both music analysis and generation, as evidenced by recent studies [4, 26, 27].

To process musical scores encoded in ABC notation more efficiently, we implement the bar patching technique [7, 25]. Bar patching involves breaking down musical sequences into units called bar patches. Each of these units

corresponds to either a bar or an information field (such as key and meter), including a sequence of characters that represent musical symbols within that patch. Unlike the conventional character-level or token-level tokenization of ABC notation [28, 29], where individual characters or tokens are processed independently, bar patching groups multiple characters into cohesive semantic units. Typically, each patch comprises 10 or more tokens, thus effectively reducing the overall sequence length of musical scores.

### 2.2 Model Architecture

As shown in Fig. 1, the MelodyT5 framework employs an encoder-decoder architecture based on the Transformer network [30], tailored for symbolic music processing. Integrating bar patching into MelodyT5 requires the incorporation of two additional components: a linear projection layer and a character-level decoder. Consequently, the model architecture encompasses the following modules:

**Linear Projection:** This component converts each bar patch into a dense embedding. It takes a multi-hot vector as input, formed by concatenating one-hot vectors representing characters within the bar patch with the shape  $S \times V$ , where  $S$  represents the patch size (i.e., the maximum number of characters in a patch) and  $V$  represents the vocabulary size. If a patch contains fewer than  $S$  characters, it will be padded with a special token to make it a  $S$ -character patch. The vector is then mapped to a dense embedding, serving as input to the patch-level encoder or decoder.

**Patch-level Encoder:** It is responsible for generating contextualized representations to understand the input musical score by operating on the dense embeddings produced by the linear projection layer. Leveraging mechanisms like self-attention and feed-forward neural networks, it captures global dependencies within the input musical score.

**Patch-level Decoder:** Tasked with generating the dense representation of the next bar patch, the patch-level decoder utilizes contextualized representations from the encoder and patch embeddings of previously generated content. It employs cross-attention and autoregressive generation mechanisms, ensuring global coherence and continuity in the sequence of generated bar patches.

**Table 1.** The MelodyHub collection statistics include the number of instances for each task along with the corresponding data sources. The JSB Chorales dataset is augmented to 15 keys due to its small size and original data in the C key.

Data Sources	Cataloging	Generation	Harmonization	Melodization	Segmentation	Transcription	Variation
<i>ABC Notation</i> [31]	184,660	184,738	31,732	31,690	—	174,779	—
<i>FolkWiki</i> [32]	6,610	6,767	1,207	1,205	—	6,218	—
<i>JSB Chorales</i> [33]	4,980	4,980	4,950	4,950	19,125	4,980	—
<i>KernScores</i> [34]	1,731	1,776	—	—	1,275	1,754	—
<i>Meertens Tune Collections</i> [35]	16,662	16,662	—	—	16,660	16,297	—
<i>Nottingham</i> [36]	1,031	1,031	1,014	1,014	—	1,021	—
<i>OpenScore Lieder</i> [37]	1,326	1,326	—	—	—	1,255	—
<i>The Session</i> [38]	44,620	44,620	3,081	3,078	—	42,838	174,104
<i>Total</i>	261,620	261,900	41,984	41,937	37,060	249,142	174,104

**Character-level Decoder:** Operating in a step-by-step manner, the character-level decoder produces characters within the next bar patch based on the dense representation generated by the patch-level decoder. By utilizing the dense representation as a context vector, it decodes each character within the bar, focusing on local information, and sequentially reconstructs every bar patch until it completes the generation of the target musical score.

The encoder-decoder architecture with bar patching in MelodyT5 enables efficient score-to-score transformations by hierarchically modelling music at both patch and character levels, capturing global structure and local details inherent in compositions.

### 2.3 Pre-training Objective

The pre-training objective of MelodyT5 aims to optimize a unified encoder-decoder framework for processing and generating symbolic music across a variety of tasks, utilizing cross-entropy loss for next token prediction.

We consider a dataset  $D$  consisting of pairs  $(X, Y)$ , where  $X$  is an input musical score and  $Y$  is the target musical score. Each score is represented as a sequence of bar patches  $\{B_1, B_2, \dots, B_n\}$ , with each bar patch  $B_i$  further decomposed into a sequence of characters  $\{c_1, c_2, \dots, c_m\}$ . The model is trained to predict each token (i.e., character) of the target score given the input score and the previously generated tokens in an autoregressive manner.

Formally, the pre-training objective can be represented as minimizing the cross-entropy loss across all tokens in the target sequence:

$$\mathcal{L}(\theta) = - \sum_{(X,Y) \in D} \sum_{i=1}^n \sum_{j=1}^m \log P_{\theta}(c_j^i | X, B_{<i}, c_{<j}^i) \quad (1)$$

where  $c_j^i$  is the  $j$ -th character in the  $i$ -th bar patch of score  $Y$ ,  $B_{<i}$  includes all bar patches before the  $i$ -th,  $c_{<j}^i$  are characters before the  $j$ -th in the current patch, and  $P_{\theta}$  is the probability of the model, parameterized by  $\theta$ , of predicting the correct character.

This objective incorporates the fundamental principle that the vast majority of symbolic music tasks can be considered as transformations from score to score, or, in other words, from an input musical score to a target musical

score. By pre-training on this objective, MelodyT5 acquires the ability to understand and replicate a wide array of patterns and structures inherent to different music tasks, which is pivotal for its success across various applications within symbolic music processing.

## 3. DATASET

This section outlines the melody curation and task definition of the MelodyHub dataset. MelodyHub, crucial for training MelodyT5, comprises seven melody-centric tasks. This collection, sourced from sheet music datasets, includes folk songs and other non-copyrighted musical scores from various traditions and epochs.

### 3.1 Melody Curation

The MelodyHub dataset was curated using publicly available sheet music datasets and online platforms, with original formats like ABC notation, MusicXML, and Humdrum. The data curation process included several steps:

1. Entries featuring explicit copyright indicators such as “copyright” or “©” symbols were excluded.
2. All data was converted to MusicXML format for standardization and subsequently transformed into ABC notation to ensure format consistency.
3. Melodies consisting of fewer than eight bars were omitted from the dataset to maintain adequate complexity and musical richness.
4. Removal of lyrics and non-musical content (e.g., contact information of transcribers and URL links) aimed to focus solely on musical notation.
5. Leading and trailing bars of complete rest were removed from each piece.
6. Each piece underwent verification for the presence of a final barline, with addition if absent.
7. Entries were deduplicated to prevent redundancy.

By ensuring the quality and consistency of the MelodyHub dataset, these steps led to a substantial collection of 261,900 melodies with uniform formatting, making it suitable for training and evaluating symbolic music models like MelodyT5.

### 3.2 Task Definition

Following the curation of melody data, the MelodyHub dataset was segmented into seven tasks, as summarized in Table 1, presented in a score-to-score format with input-output pairs. In MelodyHub, every input data includes a task identifier (e.g., `%%harmonization`) at the outset to specify the intended task. Below are the definitions of these tasks:

**Cataloging:** This task selects melodies with music-related metadata like titles, composers, and geographical origins (e.g., `C:J.S. Bach, O:Germany`). The input data includes information fields with these attributes, while specific information is removed and the order is randomized. The output includes the corresponding metadata without the musical score.

**Generation:** Here, the input solely consists of a task identifier (i.e., `%%generation`), while the output comprises comprehensive musical scores. Following TunesFormer [25], control codes are affixed to all melodies as information fields to denote musical structure information. These codes, namely `S:`, `B:`, and `E:`, signify the number of sections, bars per section, and edit distance similarity between every pair of sections within the tune.

**Harmonization:** This task involves melodies containing chord symbols. The chord symbols are removed from the input, while the original data is retained as the output. An additional information field denoting edit distance similarity (`E:`) is appended to the output, indicating the similarity between the input and output, ranging from 0 to 10 (no match at all to exact match). Lower similarity values suggest the need for more chord symbols.

**Melodization:** In contrast to harmonization, this task operates inversely and also employs melodies containing chord symbols. The notes in the original score are replaced with rests, and adjacent rest durations are combined. The resultant score, comprising rests and chord symbols, serves as the input. Similar to harmonization, an `E:` field is added at the outset of the output, with lower values facilitating the generation of more intricate melodies.

**Segmentation:** Melodies in Humdrum format (i.e., KernScores and Meertens Tune Collections) containing curly braces indicating segmentation or voices from the JSB Chorales dataset (four-part compositions) with fermatas are chosen. These markers are transformed into breath marks. The input data omits all breath marks, while the output introduces an `E:` field at the beginning to aid the generation of breath marks, with lower values implying the need for more breath marks to be added.

**Transcription:** ABC notation is initially converted to MIDI, then reconverted back to ABC. The resultant ABC from the MIDI conversion loses substantial score information, such as distinguishing enharmonic equivalents and missing musical ornaments (e.g., trill). The MIDI-converted ABC serves as the input, while the original ABC, appended with an added `E:` field, constitutes the output. Lower `E:` values denote greater discrepancies between the transcribed and input scores, particularly due to absent repeat symbols.

**Variation:** This task centres on data from The Session, wherein each ABC notation file may contain multiple variants of the same tune. Tunes with two or more variations are selected, with every possible pair of variants utilized as both input and output. The output initiates with an `E:` field signifying the extent of disparities between the input and output scores, with lower values suggesting substantial variations in the musical scores.

Together, resulting in 1,067,747 task instances in total, these tasks include various MIR challenges from analytical to generative, providing a comprehensive resource<sup>1</sup> for developing symbolic music models like MelodyT5.

## 4. EXPERIMENTS

This section evaluates the effectiveness of MelodyT5 in symbolic music processing through a series of experiments. It outlines experimental settings, conducts ablation studies on multi-task learning impact, and compares MelodyT5 with baseline models in various tasks.

### 4.1 Settings

The experiments are structured to systematically assess the capabilities of MelodyT5 for diverse symbolic music tasks. We utilize the MelodyHub dataset, which is randomly split into 99% for training and 1% for validation.

MelodyT5 features a 9-layer patch-level encoder and decoder with shared weights, a 3-layer character-level decoder, and a hidden size of 768, amounting to 113 million parameters. This configuration processes ABC sequences up to 16,384 characters, with a 256 patch length and a 64 patch size. It employs a 128-size ASCII-based vocabulary, using characters 0-2 for special tokens (pad, bos, and eos).

The AdamW optimizer [39] is used, setting a learning rate of  $2e-4$ . The process includes a 3-epoch warmup, a constant learning rate over 32 epochs, and a batch size of 10 for each GPU, ensuring consistency in hyperparameter settings across all tasks. It took approximately 2 days to complete the pre-training using 6 RTX 3090 GPUs.

In ablation studies, we investigate the effects of multi-task learning on MelodyT5, considering three settings: 1) omitting pre-training, 2) using only the downstream task-specific data from MelodyHub, or 3) utilizing the entire MelodyHub dataset, which includes all tasks.

In terms of comparative evaluations, we select open-source models that excel in their respective domains for benchmarking. MelodyT5 is fine-tuned on identical datasets to these models, ensuring fairness in comparison. For models trained on proprietary datasets, we retrain them using accessible datasets to ensure reproducibility.

Our objective evaluation strategy includes ablation studies focused on bits-per-byte (BPB) for consistent measurement, alongside task-specific metrics for comparative evaluations. Additionally, A/B tests are conducted for the subjective evaluation against baseline models.

<sup>1</sup> <https://huggingface.co/datasets/sander-wood/melodyhub>

**Table 2.** Experimental results from ablation studies illustrate the impact of multi-task learning on diverse symbolic music tasks, evaluated through BPB (bits-per-byte) to compare performance across various pre-training settings.

Pre-training	Cataloging <i>WikiMT [40]</i>	Generation <i>Wikifonia [41]</i>	Harmonization <i>CMD [42]</i>	Melodization <i>EWLD [43]</i>	Segmentation <i>Essen [44]</i>	Transcription <i>Liederschatz [45]</i>	Variation <i>The Session [38]</i>
<i>None</i>	0.0376	1.2382	0.5680	0.7949	0.0272	1.1938	0.4932
<i>Task-Specific</i>	0.0379	0.8850	0.3393	0.6322	0.0224	0.3432	–
<i>Multi-Task</i>	<b>0.0350</b>	<b>0.8472</b>	<b>0.2925</b>	<b>0.5067</b>	<b>0.0119</b>	<b>0.2969</b>	<b>0.3949</b>

### 4.2 Ablation Studies

In our ablation studies, MelodyT5 was evaluated on the test sets of various symbolic music benchmarks. Due to the lack of a directly suitable external dataset for the variation task, we chose to evaluate using the validation set of The Session. As a result, there was no task-specific pre-training for the variation task.

The ablation studies aim to explore two aspects: 1) the overall efficacy of pre-training, particularly in the context of multi-task pre-training versus task-specific pre-training, and 2) the extent to which performance gains from multi-task pre-training vary among different tasks, especially considering differences in the available volume of pre-training data across these tasks.

The ablation studies, as depicted in Table 2, show that pre-training is crucial for improving the performance of symbolic music tasks. Models trained with pre-training consistently outperform those without, indicating that pre-training enhances model generalization and performance. Multi-task pre-training is also superior to task-specific pre-training, as models trained with multi-task pre-training show lower BPB scores. This highlights the importance of leveraging multi-task pre-training to effectively capture shared patterns and structures in symbolic music data, enabling MelodyT5 to generalize better to downstream tasks.

Furthermore, it is noteworthy that while multi-task pre-training consistently yields performance gains across most tasks, the extent of improvement varies, which significantly correlates with the volume of task-specific data available for pre-training. Specifically, tasks with less data, such as segmentation and melodization, showcase more substantial performance gains from multi-task learning. On the other hand, tasks with more data, like generation and cataloging, though still benefiting from multi-task pre-training, show relatively smaller improvements. This observation suggests that while multi-task learning enhances model performance across the board, its impact is especially notable in data-constrained scenarios.

In summary, the ablation studies demonstrate the effectiveness of multi-task learning and underscore the impact of data volume on the benefits derived from such an approach. Multi-task learning boosts model performance across symbolic music tasks and provides notable advantages for tasks with limited data by leveraging shared knowledge across tasks.

### 4.3 Comparative Evaluations

For comparative evaluations, we compare MelodyT5, which is multi-task pre-trained on MelodyHub, with several task-specific baseline models, focusing on melody generation, harmonization, melodization, and segmentation. These tasks are well-established and have open-source models as competitive baselines. The following baseline models have been selected for comparison:

- **TunesFormer** [25] is applied for melody generation, featuring a Transformer-based architecture with bar patching and control codes. This approach aims to refine the efficiency of the generation process and ensure adherence to musical forms.
- **STHarm** [46] is utilized as the baseline in melody harmonization, employing a Transformer framework to convert melodies into chords. Its primary focus is on creating harmonies that preserve the structural integrity of the original melody.
- **CMT** [9] is chosen for melodization, which involves generating melodies based on chord progressions. It employs a phased training approach, conditioning the generation of rhythm and pitch on the chords to produce dynamic and coherent musical outputs.
- **Bi-LSTM-CRF** [8] is used for melody segmentation, integrating Bi-LSTM and CRF to effectively identify and segment melodic phrases for music structure analysis.

For an objective and quantifiable performance assessment that ensures reproducibility, we leverage previously established task-specific metrics. The selected metrics for our assessment include:

- **CTRL (Controllability)** [25]: Evaluates the precision of generation control through edit distance similarity between intended and actual control codes.
- **CTnCTR & PCS & MCTD** [47]: These chord/melody harmonicity metrics evaluate harmonization and melodization tasks by assessing harmonic and melodic compatibility between melodies and chords.
- **F1 Score**: Measures the balance between precision and recall in identifying correctly segmented melodic phrases.

**Table 3.** Comparative objective evaluation of the MelodyT5 model against task-specific baselines across various symbolic music tasks, utilizing task-related metrics previously established. The baselines include TunesFormer [25] for generation, STHarm [46] for harmonization, CMT [9] for melodization, and Bi-LSTM-CRF [8] for segmentation.

Models	Generation	Harmonization			Melodization			Segmentation
	CTRL $\uparrow$	CTnCTR $\uparrow$	PCS $\uparrow$	MCTD $\downarrow$	CTnCTR $\uparrow$	PCS $\uparrow$	MCTD $\downarrow$	F1 Score $\uparrow$
<i>MelodyT5</i>	<b>0.8664</b>	<b>0.7108</b>	<b>0.3274</b>	<b>1.2080</b>	0.8438	<b>0.5084</b>	<b>1.0320</b>	<b>0.9055</b>
<i>Baselines</i>	0.8162	0.5963	0.2343	1.3125	<b>0.8607</b>	0.4863	1.0610	0.8400


Table 3 shows that MelodyT5 outperforms task-specific baselines in all tasks. It surpasses the specialized baseline TunesFormer in melody generation, demonstrating enhanced control and precision in generating melodies according to specific musical forms. MelodyT5 leads in harmonization, producing chords that are harmonically compatible with the given melodies while maintaining structural coherence. Although slightly trailing CMT in CTnCTR, it still shows robust performance in other metrics for melodization, demonstrating its ability to generate melodies well integrated with chord progressions. Its performance in melody segmentation is significant, indicating its ability to accurately discern and segment melodic phrases. This performance, achieved without task-specific modifications, highlights the effectiveness of multi-task transfer learning combined with unified score-to-score transformations in symbolic music processing.

In addition to the objective metrics presented in Table 3, we recognize the limitations of solely relying on such measures to evaluate the quality of generated music. Thus, we further explored these areas (i.e., generation, harmonization, and melodization) through subjective experiments to capture listener preferences.

For our subjective evaluation, we randomly chose 30 pieces from the test set for each task and conducted blind A/B testing. Participants were presented with one randomly chosen pair from each of these 30 pairs to compare musical scores generated by MelodyT5 and baseline models under identical conditions. They were asked to choose between MelodyT5, the baseline, or no preference. Each comparison was included in two videos, showcasing both the audio and the Sibelius-rendered musical scores.

For generation, we compared the quality of melodies generated by MelodyT5 and TunesFormer, given the same control codes and information fields. In harmonization and melodization, comparisons were made against baselines given identical melodies or chords, respectively. To ensure fairness, especially considering the baseline model for melodization was limited to generating outputs of only 8 bars, we trimmed the MelodyT5-generated scores to match the output length of this baseline model.

The study involved 155 responses from students and educators with music specializations, ensuring deep understanding of melody and harmony. To secure data reliability, submissions were filtered out of those completed in less than half the overall average duration of 4 minutes and 39 seconds, i.e., those under 2 minutes and 20 seconds. This resulted in a final tally of 124 valid questionnaires.



**Figure 2.** Comparative subjective evaluation of MelodyT5 against task-specific baselines in symbolic music tasks, showing vote counts for each model.

Based on the subjective evaluation in Fig. 2, we observe a notable preference for the MelodyT5 model over the baseline in the tasks of melody generation and harmonization, with MelodyT5 receiving a higher number of votes. However, the preferences reverse in the task of melodization, where the baseline model receives a greater number of votes compared to MelodyT5. This indicates that the baseline model CMT, which employs a two-phase training process focusing separately on rhythm and pitch conditioned on chord progressions, may align more closely with human rhythmic tendencies in melodization, leading to a preference for its outputs in the subjective evaluation.

Overall, MelodyT5 excels in symbolic music processing, outperforming task-specific models in most tasks and demonstrating the effectiveness of multi-task transfer learning in this domain, despite occasional shortcomings.

## 5. CONCLUSIONS

This study presents MelodyT5, a model addressing challenges in symbolic music processing by providing a unified framework for diverse tasks. By treating music tasks as score-to-score transformations, MelodyT5 significantly improves symbolic music processing through multi-task transfer learning. Objective and subjective evaluations demonstrate that MelodyT5 generally outperforms or matches task-specific baseline models without modification. The MelodyHub dataset, with over one million task instances, offers a rich resource for training and evaluating models. While excelling in melody-centric tasks, further optimization is required to tackle more complex musical compositions, such as polyphonic arrangements.

## 6. ETHICS STATEMENT

In our research, we are committed to upholding ethical standards regarding data privacy and copyright protection. We diligently adhere to these principles throughout our data collection and processing procedures. All data utilized in our study is sourced from publicly accessible repositories, and we have taken measures to exclude any known copyrighted materials.

## 7. ACKNOWLEDGEMENTS

We would like to express our sincere gratitude to Associate Professor Bob L. T. Sturm from KTH Royal Institute of Technology for his valuable discussions and insights into the early stages of the project. We also acknowledge Yuanliang Dong and Jiafeng Liu from the Central Conservatory of Music for their assistance, and Leqi Peng from Fuyin Technology for her support with subjective experiments.

This work was supported by the following funding sources: Special Program of National Natural Science Foundation of China (Grant No. T2341003), Advanced Discipline Construction Project of Beijing Universities, Major Program of National Social Science Fund of China (Grant No. 21ZD19), and the National Culture and Tourism Technological Innovation Engineering Project (Research and Application of 3D Music).

## 8. REFERENCES

- [1] J. Liu, Y. Dong, Z. Cheng, X. Zhang, X. Li, F. Yu, and M. Sun, "Symphony generation with permutation invariant language model," in *Proceedings of the 23rd International Society for Music Information Retrieval Conference, ISMIR 2022, Bengaluru, India, December 4-8, 2022*, 2022, pp. 551–558. [Online]. Available: <https://archives.ismir.net/ismir2022/paper/000066.pdf>
- [2] P. Lu, X. Tan, B. Yu, T. Qin, S. Zhao, and T. Liu, "Meloform: Generating melody with musical form based on expert systems and neural networks," in *Proceedings of the 23rd International Society for Music Information Retrieval Conference, ISMIR 2022, Bengaluru, India, December 4-8, 2022*, 2022, pp. 567–574. [Online]. Available: <https://archives.ismir.net/ismir2022/paper/000068.pdf>
- [3] L. Min, J. Jiang, G. Xia, and J. Zhao, "Polyffusion: A diffusion model for polyphonic score generation with internal and external controls," in *Proceedings of the 24th International Society for Music Information Retrieval Conference, ISMIR 2023, Milan, Italy, November 5-9, 2023*, 2023, pp. 231–238. [Online]. Available: <https://doi.org/10.5281/zenodo.10265265>
- [4] S. Wu and M. Sun, "Exploring the efficacy of pre-trained checkpoints in text-to-music generation task," in *The AAAI-23 Workshop on Creative AI Across Modalities*, 2023. [Online]. Available: <https://openreview.net/forum?id=QmWXskBhesn>
- [5] M. Zeng, X. Tan, R. Wang, Z. Ju, T. Qin, and T. Liu, "Musicbert: Symbolic music understanding with large-scale pre-training," in *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, ser. Findings of ACL, vol. ACL/IJCNLP 2021. Association for Computational Linguistics, 2021, pp. 791–800. [Online]. Available: <https://doi.org/10.18653/v1/2021.findings-acl.70>
- [6] Z. Wang and G. Xia, "Musebert: Pre-training music representation for music understanding and controllable generation," in *Proceedings of the 22nd International Society for Music Information Retrieval Conference, ISMIR 2021, Online, November 7-12, 2021*, 2021, pp. 722–729. [Online]. Available: <https://archives.ismir.net/ismir2021/paper/000090.pdf>
- [7] S. Wu, D. Yu, X. Tan, and M. Sun, "Clamp: Contrastive language-music pre-training for cross-modal symbolic music information retrieval," in *Proceedings of the 24th International Society for Music Information Retrieval Conference, ISMIR 2023, Milan, Italy, November 5-9, 2023*, 2023, pp. 157–165. [Online]. Available: <https://doi.org/10.5281/zenodo.10265247>
- [8] Y. Zhang and G. Xia, "Symbolic melody phrase segmentation using neural network with conditional random field," in *Proceedings of the 8th Conference on Sound and Music Technology: Selected Papers from CSMT*. Springer, 2021, pp. 55–65.
- [9] K. Choi, J. Park, W. Heo, S. Jeon, and J. Park, "Chord conditioned melody generation with transformer based decoders," *IEEE Access*, vol. 9, pp. 42 071–42 080, 2021. [Online]. Available: <https://doi.org/10.1109/ACCESS.2021.3065831>
- [10] S. Wu, X. Li, and M. Sun, "Chord-conditioned melody harmonization with controllable harmonic," in *IEEE International Conference on Acoustics, Speech and Signal Processing ICASSP 2023, Rhodes Island, Greece, June 4-10, 2023*. IEEE, 2023, pp. 1–5. [Online]. Available: <https://doi.org/10.1109/ICASSP49357.2023.10096398>
- [11] S. Wu, Y. Yang, Z. Wang, X. Li, and M. Sun, "Generating chord progression from melody with flexible harmonic rhythm and controllable harmonic density," *EURASIP J. Audio Speech Music. Process.*, vol. 2024, no. 1, p. 4, 2024. [Online]. Available: <https://doi.org/10.1186/s13636-023-00314-6>
- [12] Z. Wang, K. Chen, J. Jiang, Y. Zhang, M. Xu, S. Dai, and G. Xia, "POP909: A pop-song dataset for music arrangement generation," in *Proceedings of the 21th International Society for Music Information Retrieval Conference, ISMIR 2020, Montreal, Canada, October 11-16, 2020*, 2020, pp. 38–45. [Online]. Available: <http://archives.ismir.net/ismir2020/paper/000089.pdf>

- [13] Y. Hsiao, T. Hung, T. Chen, and L. Su, “Bps-motif: A dataset for repeated pattern discovery of polyphonic symbolic music,” in *Proceedings of the 24th International Society for Music Information Retrieval Conference, ISMIR 2023, Milan, Italy, November 5-9, 2023*, 2023, pp. 281–288. [Online]. Available: <https://doi.org/10.5281/zenodo.10265277>
- [14] Y. Zhang, Z. Zhou, X. Li, F. Yu, and M. Sun, “Ccom-huqin: An annotated multimodal chinese fiddle performance dataset,” *Trans. Int. Soc. Music. Inf. Retr.*, vol. 6, no. 1, pp. 60–74, 2023. [Online]. Available: <https://doi.org/10.5334/tismir.146>
- [15] A. Holzapfel, B. L. Sturm, and M. Coeckelbergh, “Ethical dimensions of music information retrieval technology,” *Trans. Int. Soc. Music. Inf. Retr.*, vol. 1, no. 1, pp. 44–55, 2018. [Online]. Available: <https://doi.org/10.5334/tismir.13>
- [16] J. Howard and S. Ruder, “Universal language model fine-tuning for text classification,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*. Association for Computational Linguistics, 2018, pp. 328–339. [Online]. Available: <https://aclanthology.org/P18-1031/>
- [17] S. Ruder, M. E. Peters, S. Swayamdipta, and T. Wolf, “Transfer learning in natural language processing,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2, 2019, Tutorial Abstracts*. Association for Computational Linguistics, 2019, pp. 15–18. [Online]. Available: <https://doi.org/10.18653/v1/n19-5004>
- [18] M. Artetxe and H. Schwenk, “Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond,” *Trans. Assoc. Comput. Linguistics*, vol. 7, pp. 597–610, 2019. [Online]. Available: [https://doi.org/10.1162/tacl\\_a\\_00288](https://doi.org/10.1162/tacl_a_00288)
- [19] X. Liu, P. He, W. Chen, and J. Gao, “Multi-task deep neural networks for natural language understanding,” in *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*. Association for Computational Linguistics, 2019, pp. 4487–4496. [Online]. Available: <https://doi.org/10.18653/v1/p19-1441>
- [20] K. Song, X. Tan, T. Qin, J. Lu, and T. Liu, “MASS: masked sequence to sequence pre-training for language generation,” in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, ser. Proceedings of Machine Learning Research, vol. 97. PMLR, 2019, pp. 5926–5936. [Online]. Available: <http://proceedings.mlr.press/v97/song19d.html>
- [21] Z. Zhang, W. Yu, M. Yu, Z. Guo, and M. Jiang, “A survey of multi-task learning in natural language processing: Regarding task relatedness and training methods,” in *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*. Association for Computational Linguistics, 2023, pp. 943–956. [Online]. Available: <https://doi.org/10.18653/v1/2023.eacl-main.66>
- [22] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever *et al.*, “Improving language understanding by generative pre-training,” 2018.
- [23] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, 2019, pp. 4171–4186. [Online]. Available: <https://doi.org/10.18653/v1/n19-1423>
- [24] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *J. Mach. Learn. Res.*, vol. 21, pp. 140:1–140:67, 2020. [Online]. Available: <http://jmlr.org/papers/v21/20-074.html>
- [25] S. Wu, X. Li, F. Yu, and M. Sun, “Tunesformer: Forming irish tunes with control codes by bar patching,” in *Proceedings of the 2nd Workshop on Human-Centric Music Information Retrieval 2023 co-located with the 24th International Society for Music Information Retrieval Conference (ISMIR 2023), Milan, Italy, November 10, 2023*, ser. CEUR Workshop Proceedings, vol. 3528. CEUR-WS.org, 2023. [Online]. Available: <https://ceur-ws.org/Vol-3528/paper1.pdf>
- [26] L. Casini, N. Jonason, and B. L. Sturm, “Generating folk-like music in abc-notation with masked language models,” in *Ismir 2023 Hybrid Conference*, 2023.
- [27] R. Yuan, H. Lin, Y. Wang, Z. Tian, S. Wu, T. Shen, G. Zhang, Y. Wu, C. Liu, Z. Zhou, Z. Ma, L. Xue, Z. Wang, Q. Liu, T. Zheng, Y. Li, Y. Ma, Y. Liang, X. Chi, R. Liu, Z. Wang, P. Li, J. Wu, C. Lin, Q. Liu, T. Jiang, W. Huang, W. Chen, E. Benetos, J. Fu, G. Xia, R. B. Dannenberg, W. Xue, S. Kang, and Y. Guo, “Chatmusician: Understanding and generating music intrinsically with LLM,” *CoRR*, vol. abs/2402.16153, 2024. [Online]. Available: <https://doi.org/10.48550/arXiv.2402.16153>
- [28] B. L. Sturm, J. F. Santos, O. Ben-Tal, and I. Korshunova, “Music transcription modelling and composition using deep learning,” *CoRR*, vol. abs/1604.08723, 2016. [Online]. Available: <http://arxiv.org/abs/1604.08723>



- [29] C. Geerlings and A. Meroño-Peñuela, “Interacting with gpt-2 to generate controlled and believable musical sequences in abc notation,” in *NLP4MUSA*, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:227217204>
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, 2017*, pp. 5998–6008. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- [31] “Abc notation,” <http://abcnotation.com/>, accessed: 2024-04-12.
- [32] “Folkwiki,” <http://www.folkwiki.se/>, accessed: 2024-04-12.
- [33] “Chord-conditioned melody harmonization with controllable harmonicity [icassp 2023],” <https://github.com/sander-wood/deepchoir>, accessed: 2024-04-12.
- [34] “Kernscores,” <http://kern.ccarh.org/>, accessed: 2024-04-12.
- [35] “The meertens tune collections,” <https://www.liederenbank.nl/mtc/>, accessed: 2024-04-12.
- [36] “The nottingham music database,” <https://ifdo.ca/~seymour/nottingham/nottingham.html>, accessed: 2024-04-12.
- [37] “Openscore lieder corpus,” <https://musescore.com/openscore-lieder-corpus>, accessed: 2024-04-12.
- [38] “The session,” <https://thesession.org/>, accessed: 2024-04-12.
- [39] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [Online]. Available: <https://openreview.net/forum?id=Bkg6RiCqY7>
- [40] “wikimusictext Dataset on Hugging Face Datasets,” <https://huggingface.co/datasets/sander-wood/wikimusictext>, accessed: 2024-04-01.
- [41] “Download for Wikifonia all 6,675 Lead Sheets - Synth Zone Forum,” [http://www.synthzone.com/forum/ubbthreads.php/topics/384909/Download\\_for\\_Wikifonia\\_all\\_6,6](http://www.synthzone.com/forum/ubbthreads.php/topics/384909/Download_for_Wikifonia_all_6,6), accessed: 2024-04-01.
- [42] “chord-melody-dataset on GitHub,” <https://github.com/shiehn/chord-melody-dataset>, accessed: 2024-04-01.
- [43] “OpenEWLD on GitHub,” <https://github.com/00sapo/OpenEWLD>, accessed: 2024-04-01.
- [44] “KernScores: Essen Folksong Collection,” <http://kern.ccarh.org/cgi-bin/ksbrowse?l=/essen>, accessed: 2024-04-01.
- [45] “KernScores: Erk’s Liederschatz,” <https://kern.humdrum.org/cgi-bin/browse?l=users/craig/songs/erk/liederschatz>, accessed: 2024-04-01.
- [46] S. Rhyu, H. Choi, S. Kim, and K. Lee, “Translating melody to chord: Structured and flexible harmonization of melody with transformer,” *IEEE Access*, vol. 10, pp. 28 261–28 273, 2022. [Online]. Available: <https://doi.org/10.1109/ACCESS.2022.3155467>
- [47] Y. Yeh, W. Hsiao, S. Fukayama, T. Kitahara, B. Genchel, H. Liu, H. Dong, Y. Chen, T. Leong, and Y. Yang, “Automatic melody harmonization with triad chords: A comparative study,” *CoRR*, vol. abs/2001.02360, 2020. [Online]. Available: <http://arxiv.org/abs/2001.02360>