



HAL
open science

Structure-informed Positional Encoding for Music Generation

Manvi Agarwal, Changhong Wang, Gaël Richard

► **To cite this version:**

Manvi Agarwal, Changhong Wang, Gaël Richard. Structure-informed Positional Encoding for Music Generation. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2024. hal-04432659v1

HAL Id: hal-04432659

<https://hal.science/hal-04432659v1>

Submitted on 15 Feb 2024 (v1), last revised 28 Feb 2024 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

STRUCTURE-INFORMED POSITIONAL ENCODING FOR MUSIC GENERATION

Manvi Agarwal Changhong Wang Gaël Richard

LTCI, Télécom Paris, Institut Polytechnique de Paris, France

ABSTRACT

Music generated by deep learning methods often suffers from a lack of coherence and long-term organization. Yet, multi-scale hierarchical structure is a distinctive feature of music signals. To leverage this information, we propose a *structure-informed positional encoding* framework for music generation with Transformers. We design three variants in terms of absolute, relative and non-stationary positional information. We comprehensively test them on two symbolic music generation tasks: next-timestep prediction and accompaniment generation. As a comparison, we choose multiple baselines from the literature and demonstrate the merits of our methods using several musically-motivated evaluation metrics. In particular, our methods improve the melodic and structural consistency of the generated pieces.

Index Terms— symbolic music generation, Transformers, music structure, positional encoding

1. INTRODUCTION

Music generation using deep learning is seeing growing interest due to not only its commercial opportunities but also the challenges of recreating creativity with machines. Recent improvements in data-driven music generation systems have been influenced by advances in Natural Language Processing (NLP) [1], such as the use of Transformers [2] and tokenization for converting music into symbolic sequences, typically in MIDI or MIDI-inspired formats [3, 4]. Despite these advances, symbolic music generated by Transformers lacks the rich, multi-scale structures that are a characteristic feature of real music [5, 6]. In addition to their role in the enjoyment of music, such structural qualities can guide music generation by expressing knowledge about the data domain.

Previous works have hinted at the efficacy of incorporating knowledge about musical structure into data-driven models. Symbolic representations, such as REMI [3], compound

This work was funded by the European Union (ERC, HI-Audio, 101052978). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. This work was also executed with the help of HPC resources from GENCI-IDRIS (Grant 2022-AD011013986).

Code, corrected alignments and audio samples can be found on our companion website at bit.ly/structurepe

word tokens [7] and MuMIDI [4] expand the event vocabulary with tokens for musical structures, such as chord and beat. However, this obfuscates the representation of syntax or structure with that of semantics or content and causes sequence length and vocabulary size to grow. Vocabulary choice is a complex issue [8, 9] and increasing vocabulary size compromises the expressivity of multi-layered, deep models [10].

Instead, many works use the Positional Encoding (PE) module as a candidate for inserting musically-informed priors into Transformers. However, they rely on structural cues that are trivially inferred from the input. PopMAG uses bar and position [4], SymphonyNet uses note order, measure order and track ID [11] and RIPO attention uses relative pitch and onset information [12]. This leaves space to improve how structural knowledge is encoded and represented by PE.

In this paper, we explore whether providing hierarchical, musically-aware structural information, obtained non-trivially from the input by signal processing methods or human-provided annotations, to the PE module can be beneficial for Transformers for music generation.

First, we present a novel, structure-informed positional encoding framework called StructurePE. It comes in three flavours: Structure Absolute Positional Encoding (S-APE), Structure Relative Positional Encoding (S-RPE) and Structure Relative Positional Encoding with a nonstationary kernel (NS-RPE). Evaluating these variants on next-timestep prediction and accompaniment generation, we demonstrate that our methods outperform the baselines on music-focused metrics.

Second, to obtain high-quality annotations, we correct the alignment between structural labels and songs in the POP909 dataset [13], providing this via the companion website.

Third, drawing on previous work from NLP, we use Transformers without Positional Encoding (NoPE) as a baseline. We show that, though NoPE is often left out as a baseline for music generation, it is at par with other baselines on both tasks. We argue that NoPE should be included in future work on music generation.

2. METHODS

2.1. Input Representation

We use a binary pianoroll representation for the input, using a resolution of 16 timesteps for one quarter note. A pianoroll

is $\mathbf{X} \in \mathbb{B}^{(n_{\text{tracks}} \times 128) \times n_{\text{time}}}$, $\mathbb{B} = \{0, 1\}$. Every pianoroll column \mathbf{x}_t , $t \in \{1, \dots, n_{\text{time}}\}$, is associated with multiple levels of structural labels that describe the broader context of the pitches of that column. For example, \mathbf{x}_t might be part of a melody, a chord and a phrase.

2.2. Positional Encoding

Without PE, Transformers are invariant to the order of the input [2]. PE gives the model information about what content in the sequence occurs in what position. Positional information is injected into the model with two techniques: Absolute Positional Encoding (APE) and Relative Positional Encoding (RPE). As shown in Fig. 1, in APE, positional information is added to the input before it enters the Transformer. Each timestep t in the sequence is mapped to a positional index i_t , which is embedded as a vector using $p_{\text{APE}}(\cdot)$. Typically, the former is $i_t = f(t) = t$ and the latter is a sinusoidal embedding function [2]. In RPE, positional information is injected directly into the attention matrix. Each pair of timesteps (t, t') is mapped to a value that represents the relationship between them. Usually, RPE uses $f(t, t') = t - t'$. Similar to APE, RPE uses an embedding function $p_{\text{RPE}}(\cdot)$ on $f(t, t')$ [14].

2.3. Structure-informed Positional Encoding

For StructurePE, we map each timestep t to its structural vector $\mathbf{i}_t = [i_t^s]$ where $s \in S = \{\text{tempo, section, chord, mpitch}\}$. This gives us a multi-dimensional, hierarchical positional vector, each of whose elements correspond to different aspects of musical structure. We propose three variants of StructurePE.

StructureAPE (S-APE) Following the APE method [2], we use an embedding function to map each i_t^s to a vector and add these vectors to the input:

$$\text{input}_t = \mathbf{x}_t + \sum_{s \in S} p_{\text{APE}}(i_t^s) \quad (1)$$

We experimented with two embedding functions $p_{\text{APE}}(\cdot)$: a learnable embedding (L S-APE), and a sinusoidal embedding (S S-APE) [12].

StructureRPE (S-RPE) We take the RPE approach [14] and use the relative distance between i_t^s and $i_{t'}^s$ for all timesteps t, t' . The distances are embedded as vectors and incorporated directly into the attention matrix as follows:

$$z_{t,t'} = \mathbf{q}_t \mathbf{k}_{t'}^T + \sum_{s \in S} \mathbf{q}_t p_{\text{RPE}}(i_t^s - i_{t'}^s)^T \quad (2)$$

where \mathbf{q}_t is the query and $\mathbf{k}_{t'}$ is the key. Similar to S-APE, we have two choices for the embedding function $p_{\text{RPE}}(\cdot)$, giving us L S-RPE and S S-RPE.

Nonstationary StructureRPE (NS-RPE) Following a line of work that identifies the attention mechanism as a kernel [15], the standard RPE formulation, which is also used

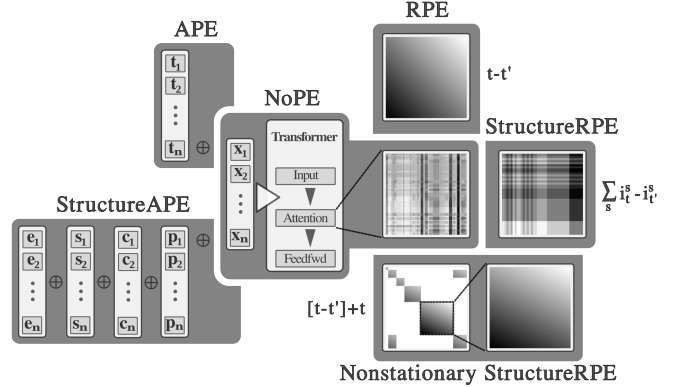


Figure 1: Illustrative schematic of PEs - both baselines (NoPE, APE and RPE) and ours (rest) - and their use in Transformers. See Sections 2.2 and 2.3 for details.

by S-RPE, can be interpreted as implementing a stationary kernel [16, 17]. S-RPE is stationary with respect to both time t and the structural positional indices i_t^s .

While the kernel formulation was introduced to reduce the complexity of the attention calculation, we can gain representational power by considering the broader category of nonstationary kernels. Nonstationary kernels can express rich relationships between positions and do not depend solely on the lag between them [18]. We add a kernel $\kappa_\varphi(t, t')$ to Eq. 2 to introduce nonstationarity with respect to time at a chosen structural level $\varphi \in S$. $\kappa_\varphi(t, t')$ gives the model information about *where* in the sequence the lag $t - t'$ occurs, creating an input-dependent kernel that is not invariant to translation.

$$z_{t,t'} = \mathbf{q}_t \mathbf{k}_{t'}^T + \sum_{s \in S} \mathbf{q}_t p_{\text{RPE}}(i_t^s - i_{t'}^s)^T + \mathbf{q}_t \kappa_\varphi(t, t')^T \quad (3)$$

$$\kappa_\varphi(t, t') = \begin{cases} p_{\text{RPE}}(t - t') + p_{\text{RPE}}(t) & \text{if } i_t^\varphi = i_{t'}^\varphi \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

To reduce computation, we only consider those pairs (t, t') which have the same structural label within the category φ (e.g., they are labelled with the same chord or same phrase). This choice is informed by our intuition that temporal relationships within the same structure should remain identical regardless of where in the sequence the structure occurs.

3. EXPERIMENTS

3.1. Task setup

We select two conditional music generation tasks:

Next-timestep prediction Given input $\{\mathbf{x}_1, \dots, \mathbf{x}_L\}$, where $\mathbf{x}_t \in \mathbb{B}^{(n_{\text{tracks}} \times 128)}$ is a column of the pianoroll for timestep t , produce as output \mathbf{x}_{L+1} .

Accompaniment generation Given input $\{\mathbf{x}_1, \dots, \mathbf{x}_L\}$, where $\mathbf{x}_t \in \mathbb{B}^{((n_{\text{tracks}}-1) \times 128)}$ gives the melody and bridge tracks, generate $\{\mathbf{y}_1, \dots, \mathbf{y}_L\}$, where $\mathbf{y}_i \in \mathbb{B}^{128}$ gives the piano track.

We vary the training length L_1 and testing length L_2 for varying levels of difficulty, parameterized as (L_1, L_2) . For next-timestep prediction, we use settings **N1**:(512, 1024) and **N2**:(1024, 1024). For accompaniment generation, we explore settings **A1**:(512, 512), **A2**:(512, 1024), and **A3**:(1024, 1024).

3.2. Model and Dataset

We use a 2-layer Transformer decoder with 4 heads, trained with a learning rate scheduler, early stopping and curriculum learning [19] on the sequence length. We use the Chinese POP909 dataset [13, 20]. Every song contains four labels, at different temporal resolutions, at each timestep - tempo, section, chord and melody - but they are not aligned with the songs due to two variable-length errors: (i) silence at the start, and (ii) musical phrases that are not part of the labeling, such as pickup measures. Although (i) is trivial to eliminate, (ii) cannot be corrected with high precision in an automated way. So, we manually select the beat position at which the structural labels should start and provide this as a supplement to the dataset on our companion website. Further details on the dataset, training procedure and data alignment method are also provided there.

3.3. Baselines

We consider two types of baselines: (i) general PE, and (ii) PE designed with structure for music generation. For (i), we have Transformers without PE (NoPE), with APE [2] and with RPE [14, 21], as presented in Section 2.2. For (ii), we consider SymphonyNet [11] as the baseline for structurally-informed APE (S-APE/b) and use note order and measure order for structural labels. We also consider RIPO attention [12] as the baseline for structurally-informed RPE (S-RPE/b) and use relative pitch and onset information as relative structural distances. We consider two possibilities for φ in NS-RPE: $\varphi = \text{chord}$ (NS-RPE/c) and $\varphi = \text{section}$ (NS-RPE/s).

Our models differ from the baselines only in the PE method being used.

3.4. Post-processing: Binarization and Velocity-encoding

To compare with the targets, we convert the probability output of the model into binary pianorolls. We adopt four methods of binarization for next-note prediction and use the mean square error to select the best method: (1) Threshold: use a fixed threshold as a step function to binarize. (2) Threshold with merge: after applying (1), fill the gaps between two occurrences of the same pitch when the gap length is below a certain value. (3) Top- k sampling: select the k highest probabilities and sample randomly from these. (4) Top- k sampling with merge: do (3), then use merge technique from (2). For accompaniment generation, we employ a velocity-encoding method to allow continuous-valued pianorolls with expressive dynamics. For this, we linearly map the probability output to loudness.

3.5. Evaluation metrics

To assess the musical quality of the generated pianorolls, we use metrics from the literature to compare the structure, melody and rhythm of the target and the prediction.

Self-similarity matrix distance (SSMD) With the chroma profile (number of occurrences of each pitch-class $C, C^\#, \dots, B$) over time for each pianoroll, we calculate the self-similarity matrix (SSM) as the cosine similarity between pairs of chroma profiles. SSMD is the mean absolute difference between the SSMs of the target and the prediction [22].

Chroma similarity (CS) The chroma onset vector of a half-measure in a pianoroll gives the number of onsets for each pitch-class in that half-measure. CS is given by the cosine similarity of chroma onset vector pairs, where one vector comes from the target pianoroll and the other comes from the corresponding half-measure of the predicted pianoroll [22].

Grooving similarity (GS) For a pianoroll, we obtain a histogram of the number of onsets, calculated at a resolution of $1/16^{\text{th}}$ notes. GS is the overlap between the note onset histograms of the target and the prediction.




Figure 2: Comparison of self-similarity matrices from music generated by baselines (columns 2 and 3) and our methods (columns 4 and 5). Best viewed in colour.

Methods	Next-timestep Prediction								Accompaniment Generation											
	N1				N2				A1				A2				A3			
	SSMD	CS	GS	NDD	SSMD	CS	GS	NDD	SSMD	CS	GS	NDD	SSMD	CS	GS	NDD	SSMD	CS	GS	NDD
	↓	↑	↑	↓	↓	↑	↑	↓	↓	↑	↑	↓	↓	↑	↑	↓	↓	↑	↑	↓
Baselines																				
NoPE	7.20	95.13	94.19	13.58	7.21	95.12	93.52	13.59	53.09	65.57	33.55	44.32	54.85	65.15	69.61	44.16	54.99	64.83	69.33	46.31
APE	10.62	88.74	91.35	21.45	6.92	93.65	90.98	12.73	49.48	68.27	22.77	43.77	54.42	66.08	48.23	43.90	49.87	68.42	45.13	43.57
RPE	6.77	93.62	92.70	12.75	6.79	93.66	90.27	12.51	49.46	66.88	30.43	48.64	50.61	66.98	65.16	55.56	49.26	68.26	48.10	61.42
S-APE/b	7.21	95.13	93.44	13.58	7.21	95.11	94.08	13.59	52.99	65.48	33.54	44.14	54.63	64.99	69.74	44.10	55.07	64.79	69.38	46.96
S-RPE/b	7.21	95.13	93.84	13.60	7.20	95.14	93.83	13.59	53.04	65.46	33.71	52.84	54.70	65.02	69.58	62.26	54.74	64.74	69.27	77.72
Our Methods																				
L S-APE	7.42	94.82	94.41	14.17	7.46	94.74	93.24	14.25	30.65	73.74	34.41	43.92	30.47	73.78	70.52	43.66	30.22	73.81	71.42	43.94
S S-APE	7.24	95.03	93.62	13.71	7.27	95.05	92.96	13.74	31.14	75.20	34.21	43.62	31.16	75.19	70.79	43.52	31.99	74.77	70.60	43.56
L S-RPE	7.20	95.05	94.84	13.58	7.19	95.07	94.62	13.59	39.27	69.73	32.82	44.60	38.23	70.82	68.92	45.25	37.56	70.73	69.30	45.74
S S-RPE	7.16	95.13	94.16	13.53	7.18	95.07	94.77	13.57	40.37	68.88	33.52	44.61	39.30	70.13	70.05	45.68	38.54	69.65	69.99	47.16
NS-RPE/c	6.84	93.76	93.17	12.84	6.85	93.83	93.19	12.77	39.99	69.99	33.22	46.37	39.17	71.24	69.17	47.43	39.64	70.88	69.44	47.63
NS-RPE/s	6.76	93.77	90.61	12.62	\	\	\	\	40.26	69.19	33.59	50.59	39.38	70.36	70.10	47.25	\	\	\	\

Table 1: Results for all tasks (see Section 3.1) and all metrics (see Section 3.5). SSMD: self-similarity matrix distance, CS: chroma similarity, GS: grooving similarity, NDD: note density distance. \uparrow : higher values are better. \downarrow : lower values are better.

Note density distance (NDD) Note density is the number of pitches in every $1/16^{th}$ note. NDD is the mean absolute difference between corresponding $1/16^{th}$ notes of the target and the prediction [23].

4. RESULTS AND DISCUSSION

We present the highlights of our findings, which are given in Table 1. The differences between the PE methods in next-timestep prediction are small. This corresponds well with our comparison of the target and generated pianorolls, which look nearly identical and sound highly similar. In contrast, the results on accompaniment generation are more varied and merit a detailed look.

We begin with SSMD since our focus is the infusion of structural information. We observe that each of our PEs outperform all baselines on SSMD in accompaniment generation, with S-APE achieving the best SSMD by a large margin. For a qualitative perspective on this finding, we plot the SSMs for one sample which falls within the 50 best SSMDs for four different PE methods on setting A2 in Fig. 2. The baseline SSMs lack even the coarse structure, while our PEs ably reproduce the large-scale similarity features. On zooming into the bottom-right section of the SSMs, outlined in red in Fig. 2, we see that S-APE misses the finer structural details. On the other hand, NS-RPE/c can capture small-scale, high-frequency information. We hypothesize that it is capable of doing so because the nonstationary kernel κ_φ from Eq. 4 models variation within uniform musical blocks. We can potentially use this to enhance the diversity of generation and produce heterogeneous structures at multiple scales.

Coming to CS, our PEs also comfortably outperform the baselines. While this result underlines the ability of our framework to model melodic features, it also indicates that our structural labels are favourable for incentivizing musically-relevant features, such as tonal consistency. The latter point is made stronger when we compare our StructurePE variants against the baseline structure-informed PEs

across tasks and metrics. The better performance of our PEs demonstrates that the type and quality of structural information is crucial.

On GS and NDD, our methods outperform the baselines but the margins are not large. This matches our observation that the generated music, both in the baselines and our methods, are indeed missing notes. However, on listening to the music generated by our models, we argue that it works better as an accompaniment, compared to the baselines. In future work, we will confirm this with a listening study.

Zooming out, we see that NoPE is at par with the other methods on most metrics for both tasks. This surprising finding agrees with previous work from NLP showing that NoPE implicitly and flexibly captures positional information [24, 25]. However, recent work on PE modules for music generation fail to include NoPE as a baseline [11, 12, 7]. We argue that NoPE should be considered a serious contender and included in future work on music generation with Transformers.

Finally, in terms of length generalization, APE performs poorly on N1, which is in line with the literature [16]. However, it performs at par (CS, SSMD) or better (NDD) than other baselines on A2. This challenges the prevailing notion that, compared to RPE, APE is bad at length generalization as it solely represents absolute positions [26].

5. CONCLUSION

In this paper, we examined how structural information about musical signals can be leveraged via PE to improve music generation with Transformers. Using a novel framework with three PE variants, we found that, compared to several baselines, our methods particularly improve the structural and melodic properties of the generated music, while also boosting rhythm and polyphony. A qualitative analysis showed that our methods reproduce the overall musical structure of the target and have the ability to represent fine structural details.

6. REFERENCES

- [1] Louis Bigo, *Modeling Musical Scores Languages*, HDR dissertation, Université de Lille, 2023.
- [2] Ashish Vaswani, Noam Shazeer, Niki Parmar, and Jakob Uszkoreit, et al., “Attention Is All You Need,” *NeurIPS*, 2017.
- [3] Yu-Siang Huang and Yi-Hsuan Yang, “Pop Music Transformer: Beat-based Modeling and Generation of Expressive Pop Piano Compositions,” *ACM International Conference on Multimedia (MM)*, 2020.
- [4] Yi Ren, Jinzheng He, Xu Tan, and Tao Qin, et al., “PopMAG: Pop Music Accompaniment Generation,” *MM*, 2020.
- [5] Shulei Ji, Xinyu Yang, and Jing Luo, “A Survey on Deep Learning for Symbolic Music Generation: Representations, Algorithms, Evaluations, and Challenges,” *ACM Comput. Surv.*, vol. 56, no. 1, 2023.
- [6] Shih-Lun Wu and Yi-Hsuan Yang, “The Jazz Transformer on the Front Line: Exploring the Shortcomings of AI-composed Music through Quantitative Measures,” *ISMIR*, 2020.
- [7] Wen-Yi Hsiao, Jen-Yu Liu, Yin-Cheng Yeh, and Yi-Hsuan Yang, “Compound Word Transformer: Learning to Compose Full-Song Music over Dynamic Directed Hypergraphs,” *AAAI*, 2021.
- [8] Wenhua Chen, Yu Su, Yilin Shen, and Zhiyu Chen, et al., “How Large a Vocabulary Does Text Classification Need? A Variational Approach to Vocabulary Selection,” *NAACL*, 2019.
- [9] Bo Zheng, Li Dong, Shaohan Huang, and Saksham Singhal, et al., “Allocating Large Vocabulary Capacity for Cross-Lingual Language Model Pre-Training,” *NAACL*, 2021.
- [10] Noam Wies, Yoav Levine, Daniel Jannai, and Amnon Shashua, “Which transformer architecture fits my data? A vocabulary bottleneck in self-attention,” *ICML*, 2021.
- [11] Jiafeng Liu, Yuanliang Dong, Zehua Cheng, and Xinran Zhang et al., “Symphony Generation with Permutation Invariant Language Model,” *ISMIR*, 2022.
- [12] Zixun Guo, J. Kang, and D. Herremans, “A Domain-Knowledge-Inspired Music Embedding Space and a Novel Attention Mechanism for Symbolic Music Modeling,” *AAAI*, 2023.
- [13] Ziyu Wang, Ke Chen, Junyan Jiang, and Yiyi Zhang, et al., “POP909: A Pop-song Dataset for Music Arrangement Generation,” *ISMIR*, 2020.
- [14] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani, “Self-Attention with Relative Position Representations,” *ACL*, 2018.
- [15] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler, “Efficient Transformers: A Survey,” *ACM Comput. Surv.*, 2022.
- [16] Antoine Liutkus, Ondřej Čífká, Shih-Lun Wu, and Umut Simsekli, et al., “Relative Positional Encoding for Transformers with Linear Complexity,” *ICML*, 2021.
- [17] Krzysztof Marcin Choromanski, Shanda Li, Valerii Likhoshesterov, and Kumar Avinava Dubey, et al., “Learning a Fourier Transform for Linear Relative Positional Encodings in Transformers,” 2023, preprint arXiv:2302.01925.
- [18] Marc G Genton, “Classes of Kernels for Machine Learning: A Statistics Perspective,” *JMLR*, vol. 2, 2001.
- [19] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston, “Curriculum learning,” *ICML*, 2009.
- [20] Shuqi Dai, Huan Zhang, and Roger B Dannenberg, “Automatic Analysis and Influence of Hierarchical Structure on Melody, Rhythm and Harmony in Popular Music,” *Joint Conference on AI Music Creativity (AIMC)*, 2020.
- [21] Cheng-Zhi Anna Huang, Ashish Vaswani, Jakob Uszkoreit, and Noam Shazeer, et al., “Music Transformer: Generating Music with Long-Term Structure,” *ICLR*, 2019.
- [22] Shih-Lun Wu and Yi-Hsuan Yang, “MuseMorphose: Full-Song and Fine-Grained Music Style Transfer with Just One Transformer VAE,” *TASLP*, 2022.
- [23] Behzad Haki, Marina Nieto, Teresa Pelinski, and Sergi Jordà Puig, “Real-time drum accompaniment using transformer architecture,” *AIMC*, 2022.
- [24] Yao-Hung Hubert Tsai, Shaojie Bai, Makoto Yamada, and Louis-Philippe Morency, et al., “Transformer Dissection: An Unified Understanding for Transformer’s Attention via the Lens of Kernel,” *EMNLP-IJCNLP*, 2019.
- [25] Adi Haviv, Ori Ram, Ofir Press, and Peter Izsak, et al., “Transformer Language Models without Positional Encodings Still Learn Positional Information,” *EMNLP*, 2022.
- [26] Amirhossein Kazemnejad, Inkit Padhi, Karthikeyan Natesan Ramamurthy, and Payel Das et al., “The Impact of Positional Encoding on Length Generalization in Transformers,” 2023, preprint arXiv:2305.19466.