

[IEEE Big Data - AIMG 2023] Paper Acceptance S44205

bigdata-inform <bigdata-inform@wi-lab.com>
Reply-To: dliao2@gmu.edu
To: hwdong@ucsd.edu, wx83@duke.edu, jmcauley@ucsd.edu, sdubnov@ucsd.edu

Sat, Nov 11, 2023 at 9:04 PM

Dear Weihan Xu, Julian McAuley, Shlomo Dubnov, and Hao-Wen Dong,

It is a pleasure to inform you that your paper (S44205), entitled

"Equipping Pretrained Unconditional Music Transformers with Instrument and Genre Controls",

has been accepted for presentation as a SHORT PAPER (5-6 pages) at the 1st Workshop on AI Music Generation with AI Music Competition. Please confirm receipt of this message.

Enclosed at the bottom of this message, please find the review report for your paper. The papers went through a rigorous review process. Each paper was reviewed by at least two members of the program committee.

NEXT STEPS - REVISION & SUBMISSION BY NOV. 19 ---

-- Please take all reviewers' comments and incorporate their suggestions into your final paper through meaningful edits. The chair and program committee will carefully examine your changes before you submit to the camera-ready.

-- It is strongly recommended that you write a brief description of your review response and any changes you make to your paper as one or two additional pages (one column, single space, 11-point font size) so that we can accurately assess your response to review comments.

-- Please submit your revision and review response as a single PDF file by the deadline listed above for the final review.

Should you have any questions, please do not hesitate to contact us. We look forward to seeing you at IEEE Big Data Conference.

Sincerely,

Workshop Chairs
1st Workshop on AI Music Generation with AI Music Competition
[https://urldefense.com/v3/__https://www.intellisky.org/workshops/workshop_AIMG2023.html__;!!Mih3wA!GEI-uO8NsCOGQxcs0bbyBW9PFi0LfXrSbwlztezS5h-FKFZevwS3mQpttSESpX3bxq5i_9GVJhVeUm4u4qr_wXmtMA\\$](https://urldefense.com/v3/__https://www.intellisky.org/workshops/workshop_AIMG2023.html__;!!Mih3wA!GEI-uO8NsCOGQxcs0bbyBW9PFi0LfXrSbwlztezS5h-FKFZevwS3mQpttSESpX3bxq5i_9GVJhVeUm4u4qr_wXmtMA$)
=====

----- Review Reports -----

The review report from reviewer #1:

*1: Is the paper relevant to the Workshop?

- No
- Yes

*2: How innovative is the paper?

- 5 (Very innovative)
- 4 (Innovative)
- 3 (Marginally)
- 2 (Not very much)
- 1 (Not)
- 0 (Not at all)

*3: How would you rate the technical quality of the paper?

- 5 (Very high)
- 4 (High)

- 3 (Good)
- 2 (Needs improvement)
- 1 (Low)
- 0 (Very low)

*4: How is the presentation?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

*5: Is the paper of interest to Workshop users and practitioners?

- 3 (Yes)
- 2 (May be)
- 1 (No)
- 0 (Not applicable)

*6: What is your confidence in your review of this paper?

- 2 (High)
- 1 (Medium)
- 0 (Low)

*7: Overall recommendation

- 5 (Strong Accept: top quality)
- 4 (Accept: a regular paper)
- 3 (Weak Accept: could be a poster)
- 2 (Weak Reject: don't like it, but won't argue to reject it)
- 1 (Reject: will argue to reject it)
- 0 (Strong Reject: hopeless)

*8: Detailed comments for the authors

The primary aim of this study is to investigate the effectiveness of the pretraining-and-finetuning paradigm in symbolic music generation. The research leverages an unprecedentedly large dataset of symbolic music, comprising 1.5 million songs, sourced from the MuseScore forum. The approach involves initially pretraining a large, unconditional transformer model on the extensive music dataset. The novelty of this research lies in the subsequent step: finetuning the pretrained model by incorporating additional control tokens. These tokens are designed to provide the model with specific controls over instrument selection and genre, a technique aimed at enhancing the model's controllability and expressiveness. The paper can be improved in perspectives of the following:

1. **Lack of Methodological Specificity:** The abstract provides an overview of using the "pretraining-and-finetuning" paradigm for symbolic music generation but falls short in detailing the specific methodologies employed. While it mentions the use of a large, unconditional transformer model and the introduction of control tokens during finetuning, it does not elaborate on the architecture of the transformer model, the nature of the control tokens, or how these tokens were integrated into the model. In the subsection 3.2 Models, the technical details are provided vaguely, which makes it hard for the audience to comprehend the overall solution of this work. This lack of detail leaves a gap in understanding the technical intricacies and innovations of the proposed approach.

2. **Experiment:** The experiment section of this paper can be improved in aspects such as more detailed discussion on the comparison between baseline methods. For example, Section 4.3 can be expanded with more detailed discussions.

=====
The review report from reviewer #2:

*1: Is the paper relevant to the Workshop?

- No
- Yes

*2: How innovative is the paper?

- 5 (Very innovative)
- 4 (Innovative)
- 3 (Marginally)
- 2 (Not very much)

- 1 (Not)
- 0 (Not at all)

*3: How would you rate the technical quality of the paper?

- 5 (Very high)
- 4 (High)
- 3 (Good)
- 2 (Needs improvement)
- 1 (Low)
- 0 (Very low)

*4: How is the presentation?

- 5 (Excellent)
- 4 (Good)
- 3 (Above average)
- 2 (Below average)
- 1 (Fair)
- 0 (Poor)

*5: Is the paper of interest to Workshop users and practitioners?

- 3 (Yes)
- 2 (May be)
- 1 (No)
- 0 (Not applicable)

*6: What is your confidence in your review of this paper?

- 2 (High)
- 1 (Medium)
- 0 (Low)

*7: Overall recommendation

- 5 (Strong Accept: top quality)
- 4 (Accept: a regular paper)
- 3 (Weak Accept: could be a poster)
- 2 (Weak Reject: don't like it, but won't argue to reject it)
- 1 (Reject: will argue to reject it)
- 0 (Strong Reject: hopeless)

*8: Detailed comments for the authors

This paper describes the interesting research on the “retrained unconditional” music transformers with instrument and genre controls. The authors collected 1.5 million sheet music files from MuseScore forum. The methods and models were described briefly. The experiments and results were reported very well. Although this paper has good writing and structure, the methods and models need to be improved. Those two parts should be described in more detail. The architecture of the system development for the experiments was not introduced either. If those points were improved and re-organized, this paper can be a formal paper, otherwise, I recommend to be a poster.

=====