

How Far Can Pretrained LLMs Go in Symbolic Music? Controlled Comparisons of Supervised and Preference-based Adaptation

Deepak Kumar*, Emmanouil Karystinaios*,
Gerhard Widmer, Markus Schedl

Johannes Kepler University Linz, Austria

deepak.kumar, emmanouil.karystinaios, gerhard.widmer, markus.schedl@jku.at

Abstract

Music often shares notable parallels with language, motivating the use of pretrained large language models (LLMs) for symbolic music understanding and generation. Despite growing interest, the practical effectiveness of adapting instruction-tuned LLMs to symbolic music remains insufficiently characterized. We present a controlled comparative study of fine-tuning strategies for ABC-based generation and understanding, comparing an off-the-shelf instruction-tuned backbone to domain-adapted variants and a music-specialized LLM baseline. Across multiple symbolic music corpora and evaluation signals, we provide some insights into adaptation choices for symbolic music applications. We highlight the domain adaptation vs. preserving prior information tradeoff as well as the distinct behaviour of metrics used to measure the domain adaptation for symbolic music.

1 Introduction

Symbolic music can be represented as discrete token sequences that exhibit long-range dependencies, hierarchical organization, and strong syntactic and structural constraints. These properties make pretrained large language models (LLMs) a great candidate for symbolic music generation and for text-conditioned musical reasoning tasks. At the same time, symbolic music differs from standard text domains in that small local errors can invalidate entire sequences, and musically meaningful structure often spans many measures, creating adaptation and evaluation challenges that are not well understood.

Recent work has demonstrated that transformer-based models can generate symbolic music and answer music-theoretic questions, either through training from scratch, music-domain pretraining, or instruction-tuned pipelines (Yuan et al., 2024;

Wang et al., 2025; Xu et al., 2025; Bhandari et al., 2025; Mundada et al., 2025). Of these works, Chat-Musician (Yuan et al., 2024) is the only work to adapt a language model for the symbolic music domain in the literature. The effectiveness of adapting general LLMs to symbolic music therefore remains insufficiently characterized. In practice, researchers and practitioners face a pragmatic question: given an off-the-shelf LLM, can it be adapted to the symbolic music domain via standard domain adaptation practices of NLP as seen for other domains (e.g., medicine, law, coding) (Ling et al., 2025; Li et al., 2024; Chen et al., 2025).

In this work, we study the adaptation of a language model to symbolic music in a controlled setting using ABC notation, a plain-text format that encodes notes, rhythms, meter, and key as ASCII characters. Its compact structure and broad public corpus availability simplify tokenization and training compared with the more verbose `**kern` or XML-based encodings. We compare a base instruction-tuned backbone to domain-adapted variants trained with supervised instruction tuning and preference-based optimization, and we include a music-specialized LLM baseline for reference. To support both symbolic generation and grounded musical reasoning, we curate and unify instruction data from multiple sources and task formats, combining composition-oriented corpora with concept-focused understanding prompts paired with ABC excerpts.

Our study addresses two research questions: i) how effectively do common finetuning strategies adapt an LLM for symbolic music understanding and generation, and ii) what tradeoffs arise between music-domain gains and retention of general-domain capability. We provide empirical comparisons across symbolic corpora and general benchmarks, and we discuss vocabulary expansion through music-specific tokens as a complementary direction for improving representation and genera-

*These authors contributed equally to this work.

tion fidelity.

2 Related Work

LLMs for symbolic music have recently been explored by representing music notation as text-like sequences, enabling next-token prediction, prompting, and instruction following over symbolic scores. Music-specialized models and pipelines have demonstrated promising capabilities for composition and music-theoretic question answering when trained or adapted on large corpora of symbolic data, often using ABC, `**kern`, or MusicXML-derived encodings, such as ChatMusician (Yuan et al., 2024), and MuPT (Qu et al., 2025).

Instruction tuning for music tasks extends these foundations by pairing natural-language prompts with symbolic outputs or analyses, with the goal of controllable generation and grounded reasoning about musical structure. Prior systems show that careful prompt templating and multi-stage generation can improve controllability, for example, by separating planning from realization or by providing intermediate constraints such as chord progressions or form descriptions, such as ComposerX (Deng et al., 2024), MuseCoco (Lu et al., 2023), etc.

Training from scratch on symbolic music, typically with music-native objectives and large-scale corpora, remains a strong alternative to adapting general LLMs. Such models can benefit from domain-specific inductive bias in data and tokenization, and may achieve strong distributional fidelity and coherence when the training scale is sufficient, e.g. NotaGen (Wang et al., 2025).

Alignment and preference learning methods, widely used in NLP to refine instruction-following behavior, have also begun to appear in music settings (Wang et al., 2025). Techniques such as preference-based optimization and related alignment objectives can be used to bias generation toward outputs that satisfy structural constraints, stylistic goals, or user intent, potentially reducing the gap between likelihood-based training and human-facing quality criteria.

3 Method

3.1 Task and Evaluation Scope

We focus on symbolic music in ABC notation and study two tasks: (1) generation or continuation,

Algorithm 1: ABC Degradation Pipeline

Input: ABC notation string A , maximum removable bars B_{\max}
Output: Degraded ABC notation string A'
 Split A into a list of lines L ;
 Initialize $L' \leftarrow L$;
Key Changing:
foreach line $l \in L'$ **do**
 if l is a key declaration **then**
 Replace key signature with a fixed alternative key (e.g., D#);
Random Pitch Swap:
 Define note pool $P = \{C, D, E, F, G, A, B, c, d, e\}$;
foreach line $l \in L'$ **do**
 if l is not a metadata line **then**
 foreach character c in l **do**
 With probability p , replace c with a random note from P if $c \in P$;
Truncate bars:
 Extract all musical (non-metadata) lines M from L' ;
if $|M| > 1$ **then**
 Sample $k \sim \mathcal{U}(1, \min(B_{\max}, |M| - 1))$;
 Randomly select k indices from M ;
 Remove the selected lines from M ;
 Set $L' \leftarrow M$;
return Concatenate lines in L' to form A'

| Statistic | Short Train | Short Test | Long Train | Long Test |
|--------------------|-------------|------------|------------|-----------|
| # Samples | 88,575 | 22,306 | 104,270 | 26,229 |
| Avg. Input Length | 178.2 | 178.4 | 201.7 | 199.8 |
| Avg. Target Length | 275.8 | 275.4 | 4,808.4 | 4,733.4 |
| Max Target Length | 500 | 500 | 615,538 | 710,425 |
| Avg. Bars | 10.2 | 10.3 | 233.0 | 229.1 |
| Notes per Bar | 6.11 | 6.14 | 4.79 | 4.77 |

Table 1: Dataset statistics for short- and long-target ABC datasets.

where the model produces ABC sequences conditioned on prompts or partial contexts, and (2) answering concept-rooted questions about music using ABC excerpts and metadata as context. In relation to both capabilities, we focus on symbolic music data from different sources, split into short sequence lengths and long sequence lengths. This setup allows us to study how LLMs adapt to symbolic music when trained on ABC data sources of varying quality and under different context length constraints.

3.2 Data Curation and Unification

We build a unified instruction-style training set by combining multiple sources that cover complementary aspects of symbolic music. We include composition-oriented corpora, such as collections of ABC tunes (monophonic) and structured symbolic datasets such as PDMX, high-quality classical music translated to ABC such as the Distant Listening Corpus (DLC) (Hentschel et al., 2025), Open Lieder (Gotham and Jonas, 2022) and Open String Quartets (Gotham et al., 2023), and understanding-

| Model | Musicpile-Short | | PDMX-Short | | ABCTunes-Short | | MMLU | Musicpile-Long | | PDMX-Long | | ABCTunes-Long | |
|--|-----------------------|----------------------|--------------------|----------------------|----------------------|------------------------|---------------------|--------------------|-----------------------|--------------------|----------------------|----------------------|-----------------------|
| | PPL↓ | FMD↓ | PPL↓ | FMD↓ | PPL↓ | FMD↓ | | ↑ | PPL↓ | FMD↓ | PPL↓ | FMD↓ | PPL↓ |
| LLaMA 3.1 Inst. 8B (Weerawardhena et al., 2025) | 693.9 _{0.0} | 66.5 _{0.2} | 9.0 _{0.0} | 320.4 _{0.1} | 28.2 _{0.0} | 467.2 _{0.3} | 0.59 _{0.0} | 6.0 _{0.0} | 130.1 _{0.0} | 3.3 _{0.0} | 475.8 _{0.0} | 5.1 _{0.0} | 449.6 _{0.1} |
| ChatMusician | 502.3 _{0.0} | 66.0 _{0.3} | 3.5 _{0.2} | 298.0 _{0.0} | 9.0 _{0.0} | 172.0 _{0.0} | 0.24 _{0.0} | 3.0 _{0.0} | 250.7 _{1.3} | 2.3 _{0.6} | na | 2.8 _{0.0} | na |
| Trained on short music sequences | | | | | | | | | | | | | |
| DPO | 401.5 _{14.2} | 72.7 _{0.0} | 9.0 _{0.4} | 349.6 _{0.0} | 517.3 _{3.9} | 348.1 _{0.0} | 0.31 _{0.0} | 6.6 _{0.2} | 233.1 _{10.5} | 3.4 _{0.3} | 502.0 _{0.0} | 21.3 _{13.4} | 240.5 _{0.0} |
| SFT | 36.3 _{0.5} | 362.1 _{1.5} | 3.7 _{0.0} | 350.3 _{0.6} | 9.3 _{0.0} | 265.25 _{23.4} | 0.29 _{0.0} | 4.5 _{0.3} | 291.3 _{0.0} | 2.4 _{0.0} | 455.5 _{5.0} | 3.4 _{0.0} | 253.9 _{7.5} |
| Trained on long music sequences | | | | | | | | | | | | | |
| DPO | 432.7 _{0.0} | 72.4 _{0.0} | 7.2 _{0.0} | 327.1 _{0.0} | 27.4 _{0.0} | 356.5 _{0.0} | 0.33 _{0.0} | 5.7 _{0.0} | 130.0 _{2.0} | 3.2 _{0.1} | 496.3 _{4.2} | 5.7 _{0.2} | 351.7 _{19.2} |
| SFT | 56.2 _{0.0} | 67.5 _{0.0} | 3.7 _{0.0} | 366.6 _{0.0} | 12.7 _{0.0} | 218.8 _{0.0} | 0.27 _{0.0} | 4.7 _{0.0} | 130.9 _{5.5} | 1.9 _{0.0} | 433.0 _{0.8} | 4.0 _{0.3} | 209.2 _{10.4} |

Table 2: Perplexity (PPL) and Fréchet Music Distance (FMD) across short sequence length symbolic music datasets and MMLU.

We report "na" for ChatMusician as results could not be calculated due to non-redable generation for FMD calculation and extremely slow evaluation due to lack padding, which restricted in resolving the issue

oriented prompts based on MusicPile-sft (Yuan et al., 2024). For the latter, we filter out examples dominated by caption-like or metadata-centric content (for example, YouTube-style descriptions) to better emphasize grounded musical concepts, terminology, and reasoning. All sources are normalized into a common ABC-centric schema with consistent separators and explicit fields for prompt, context, and target.

The data is split into two sets: long and short, based on the average sequence length of the target across the datasets. In the unified dataset, this splitting criterion is a sequence length of 500. Note that for short sequence length, the data from only three sources qualify, while for long sequence length, data from all sources qualify. Table 1 provides the key statistics of the unified datasets’ splits.

3.3 Finetuning Strategies

We compare several practical adaptation strategies that are common in NLP but underexplored for symbolic music. Our primary baseline is an off-the-shelf instruction-tuned LLM (LLaMA 3.1 Inst. 8B) (Weerawardhena et al., 2025). We then train a supervised finetuning (SFT) (Ouyang et al., 2022) variant on the unified instruction data, optimizing next-token likelihood over the concatenated prompt, context and target with standard token-level cross-entropy loss for next token prediction. In addition, we train a preference-based variant using paired outputs for the same prompt, optimizing a contrastive objective that increases preference for the chosen output relative to a rejected alternative.

For preference-based training, we construct rejected outputs via a set of musically motivated degradations applied to the chosen target, as specified in Algorithm 1, while preserving ABC syntactic validity throughout. This provides a controllable source of negative samples that differ from the cho-

sen output primarily in musical quality rather than formatting errors. We treat this signal as weak preference supervision rather than direct human preference data. Then we use the direct preference optimization (DPO) (Rafailov et al., 2023) technique to train the preference-based model.

We report results for Base, SFT, and DPO variants, and we include ChatMusician as a music-specialized reference point in our comparison. To maintain comparability with ChatMusician, we use a model from the LLaMA family (Touvron et al., 2023). We further adopt an instruction-tuned version, as instruction tuning has been shown to facilitate more effective knowledge acquisition under continued training (Jiang et al., 2024).

3.4 Metrics

We want to evaluate the LLM’s performance on different music datasets to show understanding (MusicPile) and generation (rest) capabilities. We use perplexity (PPL) to measure the token-level closeness of generation to the target for understanding and music generation. In contrast, Fréchet Music Distance (Retkowski et al., 2025) (FMD) over CLAMP2 (Wu et al., 2025) is used to measure the global-level musical similarity of the generated music. We also measure the change in prior capabilities of the LLM. In this regard, we use the Massive Multitask Language Understanding (MMLU) benchmark (Hendrycks et al., 2020) to see the cost of adapting an LLM for the music domain. MMLU is a multiple-choice question answering benchmark over 57 different subjects (e.g., history, law) with varying levels of expertise.

4 Results and Analysis

Table 2 presents the results over the short sequence length music data and the MMLU benchmark. We observe that for music understanding (i.e., Mu-

sicpile), all models show improvement compared to the base model, and the SFT variant appears to be the best-performing. For music generation, if we focus on token-level similarity (i.e., PPL), we observe ChatMusician and SFT are performing well while DPO deteriorates. In contrast, global-level similarity, as assessed using FMD, does not exhibit clear trends across datasets. Prior knowledge retention seems inversely related to global-level similarity, with DPO least affected and ChatMusician most affected. The LLM trained on long sequences also showed some improvement on the short sequence test set. However, the improvement is generally weaker than the models trained on short sequences. In contrast to the short sequence test set, Table 2 shows that the long sequence test set proves to be harder to adapt for all variants.

5 Discussion and Limitations

General Capability Retention. Finetuning for symbolic music can degrade general-domain performance, particularly when training is narrow in domain or overly aggressive in optimization. This tradeoff is especially relevant for instruction-tuned backbones that are expected to maintain broad conversational and reasoning skills. In practice, preserving general capability may require careful data mixing, conservative training budgets, and multi-objective training. The inclusion of a music-specialized baseline provides a useful reference point: music-pretrained systems can deliver strong in-domain behavior, but may deteriorate in general instruction-following and out-of-domain robustness.

Evaluation Considerations. Likelihood-based metrics such as perplexity capture token-level fit to a corpus, but they do not fully reflect musical quality, long-range structure, or instruction adherence. Distributional proxy metrics, such as FMD, can complement perplexity, but their correlation with human judgments for symbolic music is imperfect and may depend on representation and decoding settings. Overall, we highlight that metrics for symbolic music generation and understanding might be insufficient or incomplete.

Limitations. First, our preference-based training uses weak preference signals induced by musically motivated degradations while enforcing ABC syntactic validity. This reduces formatting confounds but does not replace human preference data, and

it may bias learning toward specific error patterns. Second, our study focuses on ABC as a compact text representation; results may not transfer directly to other symbolic encodings with different granularity and constraints. Third, while we curate and unify multiple data sources, coverage remains uneven across genres, instrumentation, and levels of theoretical sophistication. Fourth, evaluation of “understanding” remains sensitive to dataset construction and prompt design, and broader task coverage is left for future work. Finally, compute and context-length constraints limit exploration of very long-form structure and multi-part scores.

6 Conclusion

In this work, we presented a controlled comparative study of practical finetuning strategies for adapting instruction-tuned LLM to symbolic music in ABC notation, spanning generation and concept-rooted understanding. By comparing an off-the-shelf backbone, domain-adapted variants, and a music-specialized reference model, we characterized consistent tradeoffs between in-domain symbolic music behavior and retention of general language capabilities.

As future work, we plan to evaluate vocabulary expansion by introducing music-specific special tokens that capture recurring ABC constructs and structural markers (for example, barlines, section labels, voice separators, or compact representations of common ornaments). This requires extending the tokenizer, initializing new embeddings, and adapting the output projection head. While left for future work, we outline this procedure as a promising way to reduce fragmentation of musically meaningful units and improve generation fidelity under limited finetuning budgets. As an additional adaptation direction, we will complete vocabulary expansion experiments, expand the task suite and evaluation toward deeper music-theoretic understanding and longer-horizon generation, and investigate complementary mechanisms such as vocabulary expansion and structured prompting or planning to improve coherence, controllability, and grounded reasoning.

Acknowledgments

This work was supported by the European Research Council (ERC) under Horizon 2020 grant #101019375 “Whither Music?”, the Austrian Science Fund (FWF): Cluster of Excellence *Bilat-*

eral Artificial Intelligence (<https://doi.org/10.55776/C0E12>), the doc.funds.connect project *Human-Centered Artificial Intelligence* (<https://doi.org/10.55776/DFH23>), and the *Intent-aware Music Recommender Systems* (<https://doi.org/10.55776/P36413>). For open access purposes, the authors have applied a CC BY public copyright license to any author-accepted manuscript version arising from this submission.

References

- Keshav Bhandari, Abhinaba Roy, Kyra Wang, Geeta Puri, Simon Colton, and Dorien Herremans. 2025. Text2midi: Generating symbolic music from captions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23478–23486.
- Rubing Chen, Jiaxin Wu, Jian Wang, Xulu Zhang, Wenqi Fan, Chenghua Lin, Xiaoyong Wei, and Li Qing. 2025. Benchmarking for domain-specific llms: A case study on academia and beyond. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 11606–11619.
- Qixin Deng, Qikai Yang, Ruibin Yuan, Yipeng Huang, Yi Wang, Xubo Liu, Zeyue Tian, Jiahao Pan, Ge Zhang, Hanfeng Lin, et al. 2024. Composerx: Multi-agent symbolic music composition with llms. *arXiv preprint arXiv:2404.18081*.
- Mark R. H. Gotham, Maureen Redbond, Bruno Bower, and Peter Jonas. 2023. *The “OpenScore String Quartet” Corpus*. In *Proceedings of the 10th International Conference on Digital Libraries for Musicology*, pages 49–57, Milan Italy. ACM.
- Mark Robert Haigh Gotham and Peter Jonas. 2022. *The OpenScore Lieder Corpus*. In *Music Encoding Conference Proceedings 2021*, pages 131–136. Humanities Commons.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. In *International Conference on Learning Representations 2020*.
- Johannes Hentschel, Yannis Rammos, Markus Neuwirth, and Martin Rohrmeier. 2025. A corpus and a modular infrastructure for the empirical study of (an) notated music. *Scientific Data*, 12(1):685.
- Zhengbao Jiang, Zhiqing Sun, Weijia Shi, Pedro Rodriguez, Chunting Zhou, Graham Neubig, Xi Lin, Wen-tau Yih, and Srinu Iyer. 2024. Instruction-tuned language models are better knowledge learners. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5421–5434.
- Jiawei Li, Yizhe Yang, Yu Bai, Xiaofeng Zhou, Yinghao Li, Huashan Sun, Yuhang Liu, Xingpeng Si, Yuhao Ye, Yixiao Wu, et al. 2024. Fundamental capabilities of large language models and their applications in domain scenarios: A survey. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11116–11141.
- Chen Ling, Xujiang Zhao, Jiaying Lu, Chengyuan Deng, Can Zheng, Junxiang Wang, Tanmoy Chowdhury, Yun Li, Hejie Cui, Xuchao Zhang, et al. 2025. Domain specialization as the key to make large language models disruptive: A comprehensive survey. *ACM Computing Surveys*, 58(3):1–39.
- Peiling Lu, Xin Xu, Chenfei Kang, Botao Yu, Chengyi Xing, Xu Tan, and Jiang Bian. 2023. Musecoco: Generating symbolic music from text. *arXiv preprint arXiv:2306.00110*.
- Gagan Mundada, Yash Vishe, Amit Namburi, Xin Xu, Zachary Novack, Julian McAuley, and Junda Wu. 2025. Wildscore: Benchmarking mllms in-the-wild symbolic music reasoning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 16858–16874.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Xingwei Qu, Yinghao Ma, Ziya Zhou, Ka Man Lo, Jiaheng Liu, Ruibin Yuan, Lejun Min, Xueling Liu, Tianyu Zhang, Xeron Du, et al. 2025. Mupt: A generative symbolic music pretrained transformer. In *The Thirteenth International Conference on Learning Representations*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741.
- Jan Retkowski, Jakub Stepniak, and Mateusz Modrzejewski. 2025. Frechet music distance: A metric for generative symbolic music evaluation. In *The 39th Annual AAAI Conference on Artificial Intelligence*. AAAI (Association for the Advancement of Artificial Intelligence).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Yashan Wang, Shangda Wu, Jianhuai Hu, Xingjian Du, Yueqi Peng, Yongxin Huang, Shuai Fan, Xiaobing Li, Feng Yu, and Maosong Sun. 2025. Notagen: Advancing musicality in symbolic music generation with large language model training paradigms. *arXiv preprint arXiv:2502.18008*.

- Sajana Weerawardhena, Paul Kassianik, Blaine Nelson, Baturay Saglam, Anu Vellore, Aman Priyanshu, Supriti Vijay, Massimo Aufiero, Arthur Goldblatt, Fraser Burch, et al. 2025. Llama-3.1-foundationai-securityllm-8b-instruct technical report. *arXiv preprint arXiv:2508.01059*.
- Shangda Wu, Yashan Wang, Ruibin Yuan, Guo Zhancheng, Xu Tan, Ge Zhang, Monan Zhou, Jing Chen, Xuefeng Mu, Yuejie Gao, et al. 2025. Clamp 2: Multimodal music information retrieval across 101 languages using large language models. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 435–451.
- Weihan Xu, Julian McAuley, Taylor Berg-Kirkpatrick, Shlomo Dubnov, and Hao-Wen Dong. 2025. Generating symbolic music from natural language prompts using an llm-enhanced dataset.
- Ruibin Yuan, Hanfeng Lin, Yi Wang, Zeyue Tian, Shangda Wu, Tianhao Shen, Ge Zhang, Yuhang Wu, Cong Liu, Ziya Zhou, et al. 2024. Chatmusician: Understanding and generating music intrinsically with llm. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 6252–6271.