

Musical Score Understanding Benchmark: Evaluating Large Language Models' Comprehension of Complete Musical Scores

Congren Dai^{✉*}, Yue Yang^{✉*}, Krinos Li^{✉*}, Huichi Zhou[✉],
Shijie Liang[✉], Bo Zhang[✉], Enyang Liu[✉], Ge Jin[✉], Hongran An^{✉†}, Haosen Zhang[✉],
Peiyuan Jing[✉], Kinhei Lee[✉], Zhenxuan Zhang[✉], Xiaobing Li[✉], and Maosong Sun^{✉†‡}

[✉]Department of AI Music and Music Information Technology,
Central Conservatory of Music [✉]Imperial College London [✉]Tsinghua University

congren.dai@mail.ccom.edu.cn, imperial.ac.uk} sms@tsinghua.edu.cn

Abstract

Understanding complete musical scores entails integrated reasoning over pitch, rhythm, harmony, and large-scale structure, yet the ability of Large Language Models and Vision–Language Models to interpret full musical notation remains insufficiently examined. We introduce Musical Score Understanding Benchmark (MSU-Bench), a human-curated benchmark for score-level musical understanding across textual (ABC notation) and visual (PDF) modalities. MSU-Bench contains 1,800 generative question–answer pairs from works by Bach, Beethoven, Chopin, Debussy, and others, organised into four levels of increasing difficulty, ranging from onset information to texture and form. Evaluations of more than fifteen state-of-the-art models, in both zero-shot and fine-tuned settings, reveal pronounced modality gaps, unstable level-wise performance, and challenges in maintaining multilevel correctness. Fine-tuning substantially improves results across modalities while preserving general knowledge, positioning MSU-Bench as a robust foundation for future research in multimodal reasoning. The benchmark and code are available at <https://github.com/Congren-Dai/MSU-Bench>.

1 Introduction

Large Language Models (LLMs) and Vision–Language Models (VLMs) have recently exhibited strong capabilities in natural language understanding and generation, driving substantial advances across a broad range of Natural Language Processing tasks (Brown et al., 2020; Chowdhery et al.,

2022; OpenAI, a,b). In contrast, their ability to reason over complete musical scores remains under-explored. Existing benchmarks for musical score understanding are typically narrow in scope, concentrating on isolated fragments, short excerpts, or multiple-choice formulations, rather than supporting holistic reasoning over entire scores. Moreover, most prior work focuses on monophonic music, which consists of a single melodic line without harmonic or rhythmic accompaniment. Such settings fail to reflect the structural complexity and expressive richness required for open-ended, real-world musicological analysis.

When applied to complete scores, VLMs encounter two persistent challenges. The first is localisation: models frequently struggle to correctly identify bar positions, which is a prerequisite for answering higher-level questions related to harmony, texture, or form. For instance, when asked “Which articulation is used in bar 7?”, models often misalign the bar index and consequently return incorrect markings (see Figure 1a). The second challenge is hallucination, whereby models generate content that is not grounded in the score, often exacerbating errors introduced by incorrect localisation. Together, these issues lead to unreliable interpretations of complete scores and undermine confidence in model outputs relative to ideal, score-faithful answers (see Figure 1b).

To measure these limitations, we curate Musical Score Understanding Benchmark (MSU-Bench), a benchmark designed to evaluate the reasoning capabilities of LLMs and VLMs on complete musical scores, with particular emphasis on bar identification and higher-level musical understanding. The benchmark comprises 150 complete scores and 1,800 human-curated question–answer (QA) pairs, drawn primarily from representative textbook material used in undergraduate conservatory curricula. This educational motivation reflects the premise that models capable of answering such score-based


*Equal contribution.

†Also affiliated with the Key Laboratory of Music and Brain Science, Central Conservatory of Music, Ministry of Education, China; the Laboratory of AI Music, Central Conservatory of Music, Laboratory of Philosophy and Social Sciences, Ministry of Education, China; and Natural Language Processing and Computational Social Science Lab (THUNLP), Tsinghua University, China.

‡Corresponding author.

(a) Hallucination.

Pictures at an Exhibition
Modest Mussorgsky




Which articulation is used in bar 7?

I "saw" **staccato** markings in bar 7.

(b) Ideal scenario.

Pictures at an Exhibition
Modest Mussorgsky



Which articulation is used in bar 7?

I saw **tenuto** markings in bar 7.

Figure 1: (a) **Hallucination.** When queried about specific score features in bars, VLMs often fabricate responses that are not grounded in the actual score. (b) **Ideal scenario.** Models should accurately localise and analyse bars, thereby supporting reliable higher-level musicological reasoning.

questions could function as effective instructional assistants for music students. The benchmark is organised into four hierarchical levels of musical comprehension: Onset Information, Notation and Note, Chord and Harmony, and Texture and Form. These levels range from basic recognition of notational elements to advanced harmonic and structural analysis. By explicitly targeting localisation and hallucination while requiring reasoning over complete scores, MSU-Bench provides a systematic foundation for evaluating textual and visual musical reasoning in contemporary models.

We evaluate LLMs using textual input encoded in ABC notation and VLMs using visual input in the form of PDF representations of complete scores. Our experiments demonstrate that the aforementioned challenges can be substantially mitigated by representing scores in ABC notation (Walshaw, 2011). In this work, we treat ABC as a structured upper-bound modality for symbolic-to-theory reasoning when reliable bar and voice structure is available, while PDF inputs remain the end-to-end setting for visual score understanding. ABC notation is a text-based symbolic format that explicitly encodes bar structure, pitch, rhythm, and articulation using human-readable characters, thereby offering a compact, LLM-friendly representation; by contrast, MusicXML preserves richer notational detail but introduces longer sequences and more serialisation choices. In this work, we make the following contributions: (1) we introduce MSU-Bench, a benchmark for evaluating LLMs and

VLMs on complete musical scores, consisting of 1,800 human-curated generative QA pairs spanning four hierarchical levels of musical comprehension; (2) MSU-Bench supports multimodal evaluation through textual question answering with ABC notation and visual question answering with PDF scores; (3) zero-shot evaluation on more than fifteen models reveals a pronounced gap between textual and visual modalities, unstable performance across comprehension levels, and limited robustness; (4) fine-tuning leads to performance gains in both modalities while preserving general knowledge; and (5) we observe that posing questions jointly yields better performance than presenting them sequentially, suggesting that current models can effectively leverage hierarchical reasoning.

2 Related Work

Musical Score Representation. Musical score understanding constitutes a key task of Music Information Retrieval (MIR), aiming to analyse and interpret symbolic music representations in order to support downstream applications such as genre and style recognition (Simonetta et al., 2019). Drawing on approaches in representation learning, earlier studies have frequently employed Optical Music Recognition (OMR) to convert scores into digital formats, such as MIDI (Moore, 1988), MusicXML (Good et al., 2001), and LilyPond (Nienhuys and Nieuwenhuizen, 2003), thereby facilitating the learning of embeddings that capture musi-

(a) Raw score. Level 1: Modest Mussorgsky

Pictures at an Exhibition

♩ = 112 Promenade

Level 4: Melodic motif Level 3: G minor

(b) ABC notation's Metadata.

```
X:1
T:Pictures at an Exhibition
C:Modest Mussorgsky
L:1/4
Q:1/4=112
M:5/4
K:Bb
V:1 treble nm="Piano" snm="Pno."
V:2 bass
```

(c) Musical content in ABC notation.

```
V:1
""\n"!f!"~Promenade" !tenuto!G !tenuto!F !tenuto!B (!tenuto!c/f/ !tenuto!d) || %1
[M:6/4] (!tenuto!c/f/ !tenuto!d) !tenuto!B !tenuto!c !tenuto!G !tenuto!F || %2
[M:5/4] [B,,DG] [A,,CF] [B,,DB] ([C-A-c]/[CAf]/ [FAd]) || %3
[M:6/4] ([C-A-c]/[CAf]/ [FBd]) [DGB] [=EGc] [G,,CG] [A,,CF] || %4
[M:5/4] !tenuto!F !tenuto!G !tenuto!D (!tenuto!F/G/ !tenuto!C) || %5
V:2
z5 |[M:6/4] z6 |[M:5/4] [G,,G] [A,,F] [G,,G] [F,,F] [D,,D] || %3
[M:6/4] [F,,F] [B,,B] [G,,G] [C,,C] [=E,,=E] [F,,F] |[M:5/4] z5 || %5
```

(d) Sample questions for Levels 1, 2, 3, and 4.

Level 1: Q: Who is the composer? A: Modest Mussorgsky.	Level 2: Q: In bar 5, what is lowest note? A: Middle C (C4).	Level 3: Q: In bar 3, is there a G minor chord? A: Yes.	Level 4: Q: In which bar does the melodic motif appear? A: In bars 1 and 2.
--	--	---	---

Figure 2: Illustration of multi-level understanding in MSU-Bench using Mussorgsky’s *Pictures at an Exhibition*. Q indicates a question, while A indicates an answer. (a) Raw score excerpt with annotated tasks across four levels of difficulty. (b) Metadata encoded in ABC notation. (c) Musical content in ABC notation. (d) Sample questions for each level, showing progression from foundational concepts to higher-level musical reasoning.

cal structure and semantics for these understanding tasks (Zeng et al., 2021; Liang et al., 2020; Chou et al., 2021). On the other hand, musical notation systems such as ABC notation encode musical elements in a text format. This representation is typically written using ASCII characters (Gorn et al., 1963). Its concise, highly compressed, and language-compatible format makes it particularly suited for integration with large language models, enabling symbolic music understanding and generation (Tang et al., 2025; Wang et al., 2025a).

QA Benchmarks for Score Understanding. This area has attracted increasing interest in QA tasks for score understanding, which require more advanced forms of musical comprehension (Yue et al., 2024). Notably, MusicTheoryBench (Yuan et al., 2024) represents a systematic attempt to assess the competence of LLMs in music theory, evaluating performance across tasks that demand both music knowledge and reasoning. MusiXQA (Chen et al., 2025) evaluates VLMs in their ability to interpret musical scores represented as images. ZIQI-Eval (Li et al., 2024) benchmarks LLMs on tasks of music comprehension and generation, with particular

emphasis on their capacity to integrate contextual and cultural background knowledge. Furthermore, SSMR-Bench (Wang et al., 2025b) introduces a synthetic data generation framework capable of producing both textual and visual question formats to support comprehensive evaluations of musical understanding. Recently, WildScore (Mundada et al., 2025) introduces an in-the-wild benchmark that combines real score images with musicological multiple-choice questions (MCQs) sourced from public forums.

3 Benchmark Design

3.1 Research Questions

MSU-Bench aims to inspire future research on musical score understanding using LLMs and VLMs. In particular, it investigates the following research questions (RQs):

RQ1: How accurately can a model identify onset-level musical metadata?

Level 1 (Onset Information). Level 1 evaluates a model’s ability to extract onset-level metadata from symbolic scores, forming the basis for subsequent analysis. This includes identity information such

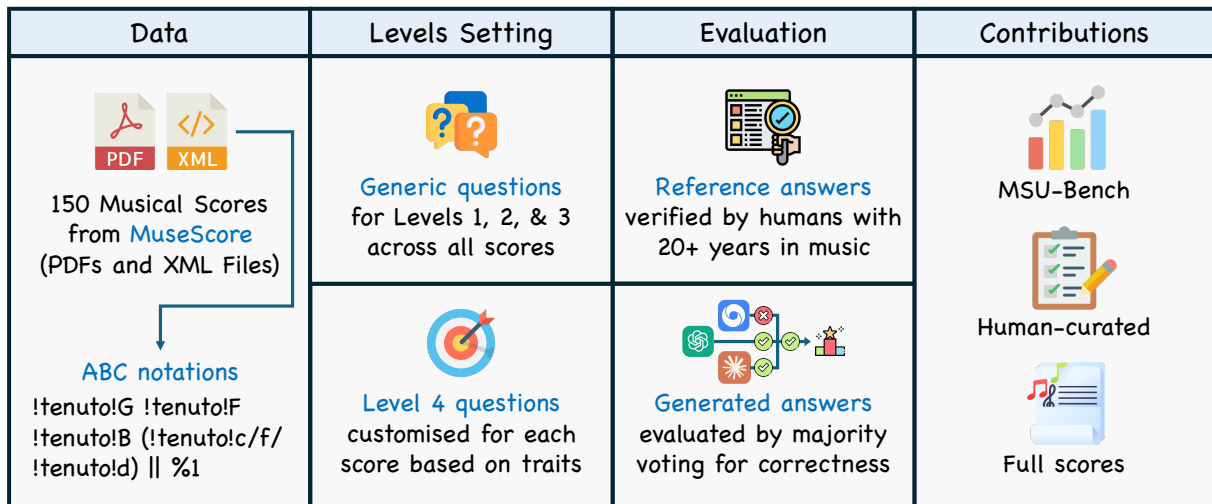


Figure 3: MSU-Bench data curation and evaluation framework. We collect 150 musical scores from MuseScore in PDF and MusicXML formats and convert them into ABC notation. Levels 1–3 use generic question templates applicable to all scores, while Level 4 questions are crafted for the unique musical features of each piece. We then verify all 1,800 questions manually. To evaluate whether generated answers match the reference answers, we employ an “LLM-as-a-judge” protocol in which LLMs independently determine whether each generated answer is correct or incorrect; the final decision is obtained through majority voting.

as composer, title, and instrumentation; notational elements such as key, clef, and time signature; performance indicators including tempo and expressive markings; and initial structural features such as the presence of an anacrusis.

RQ2: How correctly can a model interpret local notational and pitch-level features?

Level 2 (Notation & Note). Level 2 shifts from global metadata to local, bar-level notation. It assesses the identification of pitch range, accidentals, rests, ornaments, articulations, dynamics, clef, key and time signature changes, tempo changes, and repeat signs, capturing localised details essential for texture and performance understanding.

RQ3: To what extent can a model accurately analyse harmonic structures in symbolic scores?

Level 3 (Chord & Harmony). Level 3 addresses higher-order harmonic organisation, such as chord qualities and functions, inversions, voicing, spacing, and the handling of omitted or repeated notes. It further examines multi-bar progressions, cadential patterns, pedal points, ornamental harmonic devices, and the tracing of key and tonal changes from opening through modulation to tonal closure.

RQ4: To what extent can a model analyse textual and formal aspects of musical works?

Level 4 (Texture & Form). Level 4 extends harmonic analysis to global texture and form, focusing on the identification and development of melodic motifs, thematic organisation, accompa-

niment types, scoring, and orchestration. It also encompasses broader formal categories and requires sensitivity to registral distribution across bars, voices, or instruments. Level 4 prompts ask about the most frequent texture or accompaniment type in a piece, the location of principal or secondary themes, where a melodic motif first appears, how that motif is developed or ornamented, and the main structural features of the motif.

3.2 Case Study

We present a case study illustrating the structure of Level 1–4 questions in MSU-Bench, showing that ABC notation functions as a medium for musical understanding rather than a purely textual surrogate for the score, as illustrated in Figure 2a. ABC notation comprises two components, metadata (Figure 2b) and musical content (Figure 2c). Together, they encode both structural and performance-related information from Mussorgsky’s *Pictures at an Exhibition*, providing sufficient symbolic detail to support questions across all four levels (Figure 2d).

Metadata Information. The ABC header supports Level 1 questions. It specifies the title (T:Pictures at an Exhibition), composer (C:Modest Mussorgsky), default note length (L:1/4), tempo (Q:1/4=112), time signature (M:5/4), and key (K:Bb). Voice assignments are given by V:1 treble nm=“Piano” snm=“Pno.”




Figure 4: Distribution of Questions in MSU-Bench.

and V:2 bass, defining the treble and bass staves. **Musical Content.** The musical body preserves bar-level structure and all information required for Levels 2–4. In Figure 2c, bar indices are explicitly marked using %1–%5. Notation such as !tenuto!C identifies pitch-level detail within bar 5, while chord symbols [B,DG] and [G,,G,], together with the key K:Bb, imply a G-minor harmony in bar 3, where commas indicate octave displacement below the reference register. Melodic structure is encoded directly, with the motif in bar 1 given as !tenuto!G !tenuto!F !tenuto!B (!tenuto!c/f/ !tenuto!d) and its transformation in bar 2 as (!tenuto!c/f/ !tenuto!d) !tenuto!B !tenuto!c !tenuto!G !tenuto!F, preserving phrase-level information essential for higher-level musical reasoning.

3.3 Data Curation and Evaluation Process

The data collection and model evaluation processes for MSU-Bench are shown in Figure 3. The process begins with the selection of 150 scores from MuseScore. The complete list of scores included in MSU-Bench is provided in Appendix A. Our sampling unit is a score instance, which may correspond to a full score, a movement, or a self-contained excerpt, rather than a unique work title. For visual QA, the PDF of each score is employed, whereas for textual QA, the corresponding MusicXML file is converted into ABC notation. All PDFs originate from MuseScore source files and are exported directly from the corresponding project files using consistent default settings, such

as page layout and rendering parameters, avoiding scan artefacts and preserving engraving clarity. We manually inspect the exported PDFs and ensure that bar numbers are visible at the beginning of each system or segment, so bar-indexed questions are not confounded by inconsistent bar-number availability. A comprehensive set of general questions is then developed and categorised into three levels of difficulty (Levels 1–3), designed to evaluate a broad range of musical concepts encompassing fundamental notational knowledge. In addition, score-specific questions are designed for Level 4. Representative examples of Levels 1–4 questions are provided in Appendix B. With the exception of Level 1, Levels 2–4 are intentionally designed to include bar localisation tasks.

For simpler scores, we formulate questions so that each level continues to probe the intended musical ability for that piece. When a concept is not meaningfully instantiated in a particular score, we use a musicologically appropriate equivalent within the same level definition rather than reducing the task to a trivially easier variant.

To ensure high-quality reference answers, all answers are first provided manually and then verified by human domain experts possessing over 20 years of professional musical experience (detailed in Appendices C and D). This procedure guarantees that MSU-Bench is anchored in accurate and rigorously validated annotations. Each answer is carefully reviewed for correctness and completeness, and explicitly aligned with the musical content of the corresponding score.

Strict string matching is often insufficient for capturing the semantic equivalence inherent in diverse musical terminologies, especially for generative QA, such as the relationship between “V-I” and “authentic cadence”. To address this, we adopt an ensemble LLM-as-a-judge framework (Chen et al.), utilising majority voting across three distinct models, namely, ChatGPT-5 (OpenAI, b), Claude Sonnet 4 (Anthropic), and Gemini 2.5 Pro (Google), to mitigate stochasticity and bias, as illustrated in Figure 3. This methodology decouples evaluation from syntactic form, ensuring that assessment relies on semantic logic; we then verify the correctness via the human annotators. Statistical analysis confirms the reliability of this alignment, yielding a Pearson correlation (Pearson, 1896) of 0.805 ($p < 0.001$), which confirms a significant linear association. Meanwhile, a McNemar’s test (McNemar, 1947) indicates no significant difference be-

Table 1: Comparison of music-related QA benchmarks across multiple dimensions. A checkmark (✓) indicates the presence of a feature, a cross (✗) denotes its absence, and a triangle (△) represents partial coverage.

Benchmark	Modality		Sheet Music QA	Trainable	Homophony	QA Type	Quantity	Source
	Textual	Visual						
MMMU (Yue et al., 2024)	✗	✓	✓	✗	✓	MCQs	369	Web
MusiXQA (Chen et al., 2025)	✗	✓	✗	✓	✓	Generative	1.3M	Synthetic
ZIQI-Eval (Li et al., 2024)	✓	✗	✗	✗	✓	MCQs	14244	LLMs
MusicTheoryBench (Yuan et al., 2024)	✓	✗	△	✗	✗	MCQs	372	Human-labelled
SSMR-Bench (Wang et al., 2025b)	✓	✓	✓	✓	✗	MCQs	3200	Synthetic
WildScore (Mundada et al., 2025)	✓	✓	✓	✓	△	MCQs	807	In-the-wild
MSU-Bench (Ours)	✓	✓	✓	✓	✓	Generative	1800	Human-labelled

tween the LLM and human evaluations ($p > 0.05$), supporting the null hypothesis of marginal homogeneity where disagreements are random rather than systematic. In addition, the inter-judge agreement across ChatGPT-5, Claude Sonnet 4, Gemini 2.5 Pro, and an independent human assessment is 83.33%, with a 95% bootstrap confidence interval of (79.44%, 87.22%), further supporting the stability of the evaluation pipeline.

4 Benchmark Analysis

We provide a comprehensive analysis of MSU-Bench, detailing its novelty, the distribution of questions across different levels, and the characteristics of the questions.

As shown in Table 1, MSU-Bench evaluates LLMs and VLMs on complete musical scores, covering tasks from basic notation to advanced analysis. Prior benchmarks offer complementary perspectives: MMMU and ZIQI-Eval prioritise multiple-choice breadth; MusicTheoryBench provides curated questions with limited score coverage; MusiXQA scales synthetic generative queries; and SSMR-Bench targets symbolic reasoning. MSU-Bench advances this landscape by unifying textual and visual modalities, enabling model trainability, and explicitly addressing homophony in full scores, a commonly overlooked aspect. With 1,800 human-curated generative QA pairs, it balances annotation reliability with open-ended evaluation, aligning closely with current LLM and VLM research needs.

Figure 4 illustrates a balanced design in which each of the four levels accounts for 25% of MSU-Bench. More details are in Appendix E.

Level 1 emphasises performance and metadata, with onset information forming the largest proportion, complemented by smaller contributions from composer, title, tempo, and anacrusis. **Level 2** addresses markings (14%) and symbolic literacy, with

note-level features (6%), time values (3%), and ornaments (2%), and key change for modulation comprising 1%. **Level 3** distributes emphasis evenly across harmonic features, including chord identification (ID), cadences, tonic chords, and chord inversions (each 3%), with progressions, suspensions (sus.), anticipations (ant.), chord spacing, and harmonic identification (ID). **Level 4** highlights broader structural dimensions, with motif (9%) and form (8%) most prominent, supplemented by texture, register, tone, and other questions.

In addition, MSU-Bench encompasses a wide range of composers, as shown in Figure 7, Appendix F, spanning historical periods and stylistic traditions including the Baroque, Classical, Romantic, and twentieth-century repertoire. The distributions of scores by period and genre are presented in Figures 8a and 8b from Appendix F.

5 Experiments

5.1 Experiment Settings

Evaluation. Model outputs are evaluated through a voting process involving ChatGPT-5, Claude Sonnet 4, and Gemini 2.5 Pro, with more details provided in Appendices G and H. Accuracy is reported at both the individual level and the aggregate level (overall). We consider two evaluation settings: (1) zero-shot, testing models directly on the 1,800 QA pairs; and (2) fine-tuned, where models are adapted with LoRA (Hu et al., 2021). We also introduce the Level-wise Success Rate (LSR), defined as the proportion of scores for which all questions are answered correctly from Level 1 through Level l . Let n denote the maximum level, let $l \in \{1, 2, \dots, n\}$ be the level index, and let \mathcal{S} denote the score set. Then, the LSR at Level l is defined as

$$LSR(l) = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \mathbf{1} \left[\bigwedge_{j=1}^l \text{AllCorrect}(s, j) \right],$$

Table 2: Zero-shot evaluation results on MSU-Bench (full), with the highest accuracy in **bold**. We evaluate 12 questions per score in a single run for each model to report the accuracy for each level and overall.

Models	<i>Musical Score Understanding Benchmark</i>				
	Level 1 (450)	Level 2 (450)	Level 3 (450)	Level 4 (450)	Overall (1800)
<i>Textual QA</i>					
Qwen3-VL-235B-A22B-Instruct	58.44	33.56	34.00	38.89	41.22
Claude Opus 4	57.11	36.89	35.56	35.56	41.28
Claude Sonnet 4	61.11	40.67	35.56	33.11	42.61
Grok 4	62.00	40.00	31.11	37.11	42.61
ChatGPT-5-mini	59.11	43.56	31.33	40.89	43.72
ChatGPT-5	62.00	50.22	38.44	38.44	47.28
Gemini 2.5 Pro	65.33	56.00	38.67	37.78	49.44
<i>Visual QA</i>					
ChatGPT-5-mini	7.11	6.67	8.89	6.67	7.33
Grok 4	14.00	11.11	18.44	21.33	16.22
Qwen3-VL-235B-A22B-Instruct	18.67	15.33	22.44	22.67	19.78
Gemini 2.5 Flash	19.56	15.33	29.56	18.00	20.61
Qwen2.5-VL-72B-Instruct	21.78	18.22	27.33	18.89	21.56
Claude Sonnet 4	27.11	16.44	27.33	18.44	22.33
Gemini 2.5 Pro	22.00	22.44	29.11	20.00	23.39
Claude Opus 4	25.33	21.78	30.44	19.33	24.22

where $\text{AllCorrect}(s, j)$ is true if all Level- j questions for score s are answered correctly, and $\mathbf{1}[\cdot]$ is the indicator function. Then, we use the Wilson score interval (Wilson, 1927) to calculate the 95% Confidence Interval (CI) for the LSR at Level l , which is given by

$$\text{CI}(l) = \frac{\hat{p}_l + \frac{z^2}{2n_l}}{1 + \frac{z^2}{n_l}} \pm \frac{z}{1 + \frac{z^2}{n_l}} \sqrt{\frac{\hat{p}_l(1 - \hat{p}_l)}{n_l} + \frac{z^2}{4n_l^2}},$$

where \hat{p}_l is the LSR at Level l , $n_l = |S|$, and z is the standard normal quantile ($z = 1.96$).

Baselines. We evaluate a diverse set of models for the zero-shot evaluation, including both LLMs and VLMs. For textual QA (ABC notation), we evaluate ChatGPT-5, ChatGPT-5-mini (OpenAI, b), Claude Opus 4 (Anthropic), Claude Sonnet 4, Deepseek-V3 (DeepSeek-AI), Gemini 2.5 Flash (Google), Gemini 2.5 Pro, Grok 4 (xAI, 2025), Llama 4 Maverick (Meta AI), Llama 4 Scout (Meta AI), Qwen2.5-VL-7B-Instruct (Qwen Team, a), Qwen2.5-VL-32B-Instruct (Qwen Team, a), Qwen2.5-VL-72B-Instruct (Qwen Team, a), Qwen3-4B (Qwen Team, b), Qwen3-32B (Qwen Team, b), Qwen3-Max (Qwen Team, b), and Qwen3-VL-235B-A22B-Instruct (Qwen Team, c).

For visual QA (PDF documents), we include Claude Opus 4, Claude Sonnet 4, Gemini 2.5 Flash, Gemini 2.5 Pro, ChatGPT-5-mini, Grok 4, Qwen2.5-VL-32B-Instruct, Qwen2.5-VL-72B-Instruct, and Qwen3-VL-235B-A22B-Instruct.

Models. We employ Qwen3-0.6B (Qwen Team, b), Qwen3-1.7B (Qwen Team, b), Qwen3-4B, and Qwen2.5-VL-3B-Instruct for the fine-tuned evaluation, adapted using LoRA.

Data Splitting. MSU-Bench consists of 150 musical scores. It is divided into training, validation, and testing sets in a 6:2:2 ratio, corresponding to 90, 30, and 30 pieces, respectively. The 30-score test set is fixed and never used for training. For the fine-tuned evaluation, all test scores are identical to those used in the zero-shot evaluation.

Training. We fine-tune the models for 20 epochs on $6 \times \text{A800}$ GPUs using LoRA (rank 8). Training uses AdamW (Loshchilov and Hutter, 2019) with a 5×10^{-5} learning rate, cosine schedule, 10% warm-up, batch size 1, and gradient accumulation of 16. For Qwen2.5-VL-3B-Instruct, we consider three types of input: PDF only, ABC notation only, and their combination (detailed in Appendix I).

5.2 Empirical Results

Zero-shot Evaluation. In Table 2, models perform substantially better on the textual QA setting than on the visual QA setting. In textual QA, Gemini 2.5 Pro achieves the best overall accuracy (49.44%), excelling particularly at Level 1 (65.33%) and Level 2 (56.00%). ChatGPT-5 follows closely (47.28%), demonstrating strong stability on higher-level questions (Levels 3–4). Notably, ChatGPT-5-mini attains the highest accuracy on

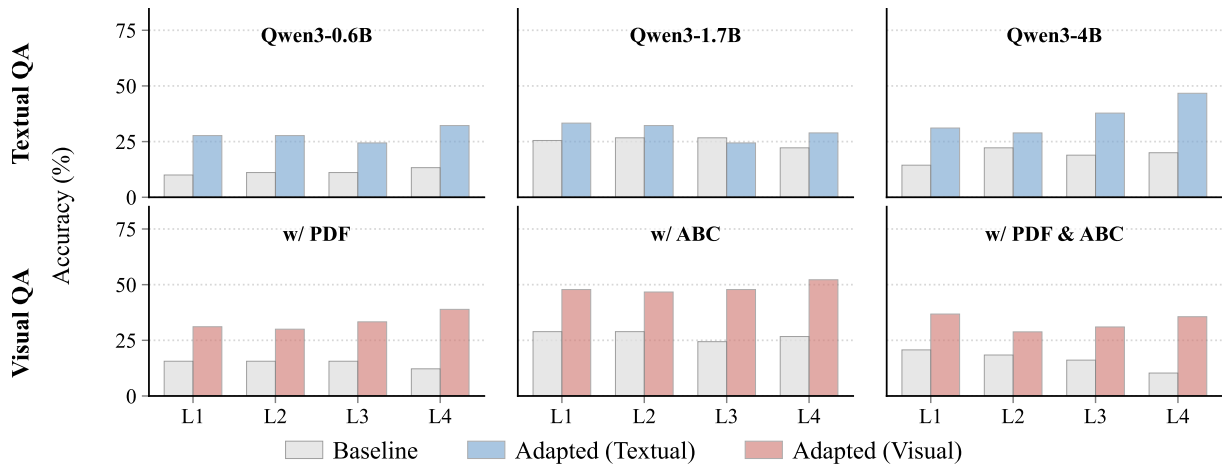


Figure 5: Performance of baseline and LoRA-adapted models on MSU-Bench (testing set). Qwen2.5-VL-3B-Instruct is adapted using LoRA across the three input modalities outlined in Section 5.1. L denotes Levels.

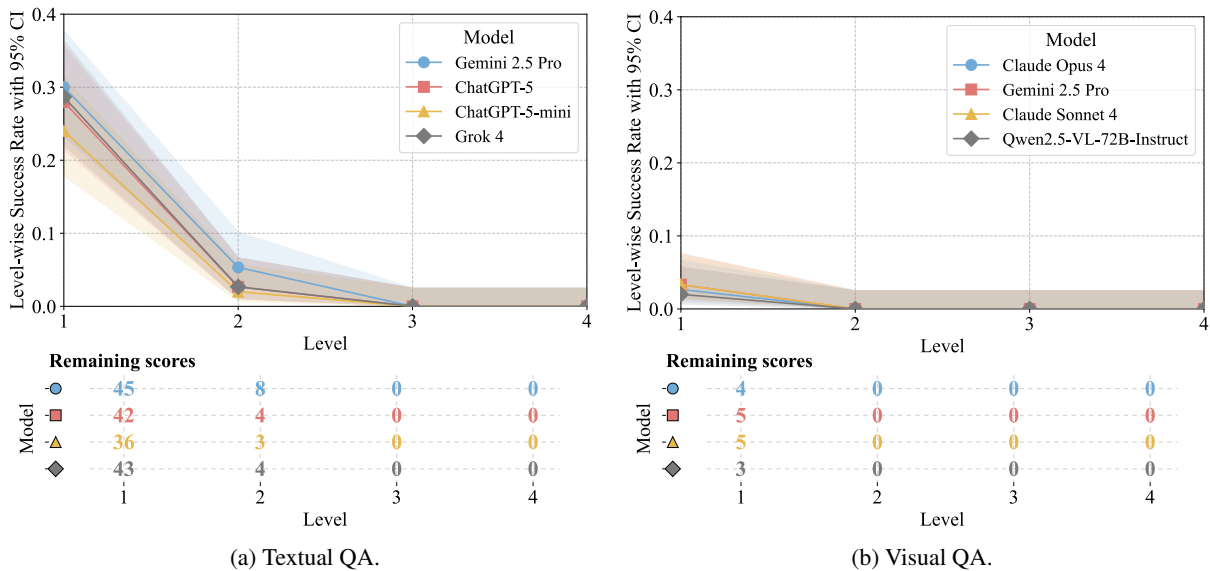


Figure 6: Level-wise Success Rate for zero-shot evaluation on MSU-Bench.

Level 4 (40.89%), suggesting an advantage in more complex reasoning despite its smaller size. Claude Opus 4, Claude Sonnet 4, Grok 4, and Qwen3-VL-235B-A22B-Instruct achieve comparable overall performance, with accuracies ranging from 41.22% to 42.61%. More results are reported in Table 5 in Appendix J. The evaluation times of models achieving more than 40% overall accuracy are reported in Appendix K (see Figure 9). While models such as Gemini 2.5 Pro, ChatGPT-5, and ChatGPT-5-mini achieve the highest levels of accuracy, their evaluation times are substantially longer (more than 11 hours). Notably, Qwen3-VL-235B-A22B-Instruct requires only approximately one hour to achieve an overall accuracy of 41.22%. For the proprietary models in Table 2, we report the corresponding

95% percentile bootstrap confidence intervals in Table 3. Notably, in textual QA, the confidence intervals for Gemini 2.5 Pro and ChatGPT-5-mini do not overlap, suggesting a meaningful separation under this CI estimate. For visual QA, overall accuracies are markedly lower (around 20%), with the strongest model (Claude Opus 4) reaching only 24.22%. It also achieves the highest Level 3 accuracy (30.44%), while Claude Sonnet 4 leads at Level 1 (27.11%) and Gemini 2.5 Pro at Level 2 (22.44%). These results highlight a clear modality gap: ABC notation provides a more reliable representation than raw score images, where recognition and localisation errors continue to dominate performance. Some PDF-QA failure cases are provided in Appendix L.

Table 3: Accuracy and 95% percentile bootstrap confidence intervals for the top-performing proprietary zero-shot models in each modality. The interval is computed as $CI_{95\%} = [Q_{0.025}(\{\hat{p}^{*(b)}\}_{b=1}^B), Q_{0.975}(\{\hat{p}^{*(b)}\}_{b=1}^B)]$, where $\hat{p}^{*(b)}$ denotes the accuracy estimate obtained from bootstrap resample b . All values are reported in percentage points.

Modality	Model	Accuracy (%)	95% CI (%)
Textual QA	Grok 4	42.61	[40.34, 44.91]
	ChatGPT-5-mini	43.72	[41.45, 46.02]
	ChatGPT-5	47.28	[44.98, 49.59]
	Gemini 2.5 Pro	49.44	[47.14, 51.75]
Visual QA	Gemini 2.5 Flash	20.61	[18.81, 22.54]
	Claude Sonnet 4	22.33	[20.47, 24.32]
	Gemini 2.5 Pro	23.39	[21.49, 25.40]
	Claude Opus 4	24.22	[22.30, 26.25]

In Figure 12 within Appendix M, we observe that the models achieve superior performance when the 12 questions pertaining to a score are posed simultaneously, particularly for Claude Sonnet 4, Claude Opus 4, and Qwen2.5-VL-72B-Instruct. This finding suggests that presenting all questions in a batch may allow Level 1 answers to support higher-level responses in Table 2.

Fine-tuned Evaluation. Figure 5 details the results following LoRA fine-tuning, which yields substantial gains across textual and visual QA. The textual QA results illustrate a distinct scaling trend within the Qwen3 family; accuracy improves as model size increases from 0.6B to 4B, with Qwen3-4B proving the most robust. Notably, performance improves as question complexity progresses from L1–L4, suggesting that LLMs possess strong logical inference abilities for solving harder questions using ABC notation.

For visual QA, Qwen2.5-VL-3B-Instruct exhibits a significant advantage when using ABC inputs only, clearly outperforming Qwen3-4B, likely because ABC reduces the visual noise inherently present in PDF files.

LSR Analysis. Figures 6a and 6b show the LSR from Levels 1–4 across all 150 scores in MSU-Bench. In textual QA (Figure 6a), models perform moderately at Level 1 (25–35%), with Gemini 2.5 Pro slightly ahead of ChatGPT-5 and Grok 4, but drop below 10% by Level 2 and nearly vanish by Level 3. Visual QA (Figure 6b) is worse: models start at 5–10% on Level 1 and collapse almost entirely by Level 2. The “remaining scores” count reinforces this gap, with about 41.5 scores surviving past Level 1 in textual QA versus only 4.25 in visual QA. These results highlight that while models can solve isolated questions, sustaining correctness across levels is extremely difficult, underscoring LSR’s diagnostic strictness; LSR for LoRA-

adapted models is also reported in Appendix N.

Massive Multitask Language Understanding (MMLU). As shown in Table 6 from Appendix O, the models adapted with LoRA exhibit minimal forgetting, with performance remaining close to that of their base versions tested on MMLU (Hendrycks et al., 2021). These results indicate that LoRA adaptation effectively preserves the models’ general knowledge while enhancing their capabilities in musical score understanding and reasoning.

Table 4: Performance evaluation using title-only input. L1–L4 denote Levels 1–4 of MSU-Bench.

Models	L1	L2	L3	L4
Qwen3-0.6B	17.7%	14.4%	11.1%	10.0%
Qwen3-1.7B	23.3%	22.2%	24.4%	24.4%
Qwen3-4B	24.4%	30.0%	26.6%	23.3%
Qwen2.5-VL-3B	15.5%	22.2%	20.0%	15.6%

Analysis of Score Content Reliance. Following Kumar et al. (2025) and Zang et al. (2025), we test whether models use local musical cues or rely on title-based guessing. In Table 4, title-only input yields weak performance across Levels 1–4, suggesting that answers cannot be reliably inferred from titles alone and that models instead depend on musical information provided in ABC notation or PDF scores when such input is available.

6 Conclusion

We introduce Musical Score Understanding Benchmark (MSU-Bench), a benchmark for evaluating complete musical score reasoning in LLMs and VLMs across textual (ABC) and visual (PDF) modalities. Our experiments reveal that while current models struggle with multi-level comprehension, particularly in visual QA, the use of ABC notation can mitigate part of this issue by providing precise bar information.

Limitations

While ABC notation effectively mitigates visual processing errors, the substantial performance gap observed between textual and visual modalities highlights the limitations of current models in interpreting raw sheet music without symbolic aid. In this work, ABC should therefore be interpreted as a structured upper-bound modality for symbolic reasoning rather than as evidence of fully solved visual score understanding. Computational constraints regarding token limits necessitated the exclusion of certain longer scores during the multimodal fine-tuning experiments. In addition, the scope of MSU-Bench is primarily centred on the Western art music canon, spanning the Baroque to the early twentieth century. We seek to broaden this musical scope to encompass non-Western, contemporary, and more diverse notational practices in future work.

Acknowledgements

This work is supported by the Advanced Discipline Construction Project of Beijing Universities, the Special Programme of National Natural Science Foundation of China (Grant No. T2341003), and the Major Programme of National Social Science Fund of China (Grant No. 21ZD19).

References

- Anthropic. 2025. [Introducing claude 4](#).
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Dongping Chen, Ruoxi Chen, Shilin Zhang, Yaochen Wang, Yinuo Liu, Huichi Zhou, Qihui Zhang, Yao Wan, Pan Zhou, and Lichao Sun. [Mllm-as-a-judge: Assessing multimodal llm-as-a-judge with vision-language benchmark](#). In *Forty-first International Conference on Machine Learning*.
- Jian Chen, Wenye Ma, Penghang Liu, Wei Wang, Tengwei Song, Ming Li, Chenguang Wang, Jiayu Qin, Ruiyi Zhang, and Changyou Chen. 2025. [Musixqa: Advancing visual music understanding in multimodal large language models](#). *Preprint*, arXiv:2506.23009.
- Yi-Hui Chou, I-Chun Chen, Chin-Jui Chang, Joann Ching, and Yi-Hsuan Yang. 2021. [Midibert-piano](#): Large-scale pre-training for symbolic music understanding. *CoRR*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, and 48 others. 2022. [Palm: Scaling language modeling with pathways](#). *Preprint*, arXiv:2204.02311.
- DeepSeek-AI. 2025. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Michael Good and 1 others. 2001. [Musicxml: An internet-friendly format for sheet music](#). In *Xml conference and expo*, pages 03–04.
- Google. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *Preprint*, arXiv:2507.06261.
- Saul Gorn, Robert W Bemer, and Julien Green. 1963. [American standard code for information interchange](#). *Communications of the ACM*, 6(8):422–426.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Sonal Kumar, Šimon Sedláček, Vaibhavi Lokegaonkar, Fernando López, Wenyi Yu, Nishit Anand, Hyeon-gon Ryu, Lichang Chen, Maxim Plička, Miroslav Hlaváček, William Fineas Ellingwood, Sathvik Udupa, Siyuan Hou, Allison Ferner, Sara Barahona, Cecilia Bolaños, Satish Rahi, Laura Herrera-Alarcón, Sathvik Dixit, and 15 others. 2025. [Mmau-pro: A challenging and comprehensive benchmark for holistic evaluation of audio general intelligence](#). *Preprint*, arXiv:2508.13992.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). In *Proceedings of the 29th ACM Symposium on Operating Systems Principles, SOSP '23*, page 611–626, New York, NY, USA. Association for Computing Machinery.
- Jiajia Li, Lu Yang, Mingni Tang, Cong Chen, Zuchao Li, Ping Wang, and Hai Zhao. 2024. [The music maestro or the musically challenged, a massive music evaluation benchmark for large language models](#). *Preprint*, arXiv:2406.15885.

- Hongru Liang, Wenqiang Lei, Paul Yaozhu Chan, Zhenglu Yang, Maosong Sun, and Tat-Seng Chua. 2020. Pirhdy: Learning pitch-, rhythm-, and dynamics-aware embeddings for symbolic music. In *Proceedings of the 28th ACM international conference on multimedia*, pages 574–582.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). *Preprint*, arXiv:1711.05101.
- Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157.
- Meta AI. 2024. [The llama 4 herd: The beginning of a new era of natively multimodal ai innovation](#).
- F Richard Moore. 1988. The dysfunctions of midi. *Computer music journal*, 12(1):19–28.
- 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](#). *Preprint*, arXiv:2509.04744.
- Han-Wen Nienhuys and Jan Nieuwenhuizen. 2003. Lilypond, a system for automated music engraving. In *Proceedings of the xiv colloquium on musical informatics (xiv cim 2003)*, volume 1, pages 167–171. Citeseer.
- OpenAI. 2024a. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- OpenAI. 2025b. [Introducing gpt-5](#).
- Karl Pearson. 1896. Mathematical contributions to the theory of evolution. iii. regression, heredity, and panmixia. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 187:253–318.
- Qwen Team. 2025a. [Qwen2.5-vl technical report](#). *Preprint*, arXiv:2502.13923.
- Qwen Team. 2025b. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Qwen Team. 2025c. [Qwen3-vl: Sharper vision, deeper thought, broader action](#). <https://qwen.ai/blog?id=99f0335c4ad9ff6153e517418d48535ab6d8afef&from=research.latest-advancements-list>. Accessed: 2025-09-24.
- Federico Simonetta, Stavros Ntalampiras, and Federico Avanzini. 2019. Multimodal music information processing and retrieval: Survey and future challenges. In *2019 international workshop on multilayer music representation and processing (MMRP)*, pages 10–18. IEEE.
- Mingni Tang, Jiajia Li, Lu Yang, Zhiqiang Zhang, Jinghao Tian, Zuchao Li, Lefei Zhang, and Ping Wang. 2025. Nota: Multimodal music notation understanding for visual large language model. *arXiv preprint arXiv:2502.14893*.
- Chris Walshaw. 2011. The ABC music standard 2.1. https://michaeleskin.com/abctools/abc_standard_v2.1.pdf. Dec 2011. Accessed: 2026-02-17.
- Yashan Wang, Shangda Wu, Jianhui Hu, Xingjian Du, Yueqi Peng, Yongxin Huang, Shuai Fan, Xiaobing Li, Feng Yu, and Maosong Sun. 2025a. Notagen: Advancing musicality in symbolic music generation with large language model training paradigms. *arXiv preprint arXiv:2502.18008*.
- Zhilin Wang, Zhe Yang, Yun Luo, Yafu Li, Haoran Zhang, Runzhe Zhan, Derek F. Wong, Jizhe Zhou, and Yu Cheng. 2025b. [Synthesizing sheet music problems for evaluation and reinforcement learning](#). *Preprint*, arXiv:2509.04059.
- Edwin B Wilson. 1927. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212.
- xAI. 2025. Grok 4. <https://x.ai/news/grok-4>. Model Card updated August 20, 2025.
- Ruibin Yuan, Hanfeng Lin, Yi Wang, Zeyue Tian, Shangda Wu, Tianhao Shen, Ge Zhang, Yuhang Wu, Cong Liu, Ziya Zhou, Ziyang Ma, Liumeng Xue, Ziyu Wang, Qin Liu, Tianyu Zheng, Yizhi Li, Yinghao Ma, Yiming Liang, Xiaowei Chi, and 16 others. 2024. [Chatmusician: Understanding and generating music intrinsically with llm](#). *Preprint*, arXiv:2402.16153.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567.
- Yongyi Zang, Sean O’Brien, Taylor Berg-Kirkpatrick, Julian McAuley, and Zachary Novack. 2025. [Are you really listening? boosting perceptual awareness in music-qa benchmarks](#). *Preprint*, arXiv:2504.00369.
- Mingliang Zeng, Xu Tan, Rui Wang, Zeqian Ju, Tao Qin, and Tie-Yan Liu. 2021. Musicbert: Symbolic music understanding with large-scale pre-training. *arXiv preprint arXiv:2106.05630*.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyang Luo. 2024. [LlamaFactory: Unified efficient fine-tuning of 100+ language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 400–410, Bangkok, Thailand. Association for Computational Linguistics.

A Appendix: List of Musical Scores

1. Cello Suite No.1 BWV 1007 - 1. Prélude
2. Solfeggietto in C minor
3. Toccata and Fugue in D minor BWV 565
4. Fugue in G Minor BWV 542
5. Fugue I in C major BWV 846
6. Fugue in D minor BWV 948
7. Fugue in G Minor BWV 578
8. Prelude I in C major BWV 846
9. Sonate No. 16 1st Movement
10. Piano Sonata No. 5 in C Minor Op.10 No.1
11. Sonata in G Op.14 No.2 Movement 1
12. Piano Sonata in A major Op.2 No.2
13. Piano Sonata No. 3 in C Major Op. 2 No. 3
14. Sonata No. 23 Op. 57 Appassionata
15. Sonata Op.31 No.17 in D minor Tempest
16. Piano Sonata No. 17 in D minor Op. 31 No. 2
17. Piano Sonata No.18 in E flat major Op.31 No.3
18. Sonate No.8 Op.13 Pathétique 3 Rondo. Allegro Sonata No.8
19. Les Nuits d'été
20. Symphonie fantastique, H 48
21. Polovtsian Dances
22. Hungarian Dance No. 5
23. Rhapsody Op. 79 No. 2
24. Waltz Op.39 No.3
25. Intermezzo in E flat major Op.117 No.1
26. B minor Rhapsody 1 Op. 79
27. Ballade Op.118 No.3
28. Intermezzo Op. 116 No. 2
29. Intermezzo Op. 118 No. 2 A Major
30. Violin Concerto in E minor Op.64
31. Lullaby Op.49 No.4
32. Waltz in A Major Op.39 No.15
33. Fantaisie-Impromptu in C♯ Minor
34. Nocturne Op. 9 No.1
35. Nocturne-No. 20 in C Sharp Minor
36. Ballade no.1 in G minor Op.23
37. Sonata No.2 Op.35 1st Movement
38. Ballade No.3 in A flat major Op.47
39. Ballade No.4 in F minor Op
40. Prélude Opus 28 No. 4 in E Minor
41. Waltz in A Minor
42. Nocturne Op. 27, No. 2
43. Suite Bergamasque
44. La fille aux cheveux de lin
45. Reverie
46. Clair de lune
47. Premier Trio
48. Syrinx
49. Sonate pour Violoncelle et Piano
50. Symphony No. 9 New World II, Largo
51. Symphony No. 9 New World:IV, Allegro con fuoco
52. Humoresque No.7
53. Holberg Suite Op.40 I.Praeludium
54. Wedding Day at Troldhaugen
55. Anitras Dance Piano solo
56. Dance Op. 12 No. 4
57. Sailors Song Op.68 No.1
58. Waltz Op. 12 No. 2
59. Butterfly Sommerfugl Op. 43 No. 1
60. Piano Concerto in A minor Op.16

- | | |
|---|--|
| 61. In the Hall of the Mountain King | 91. Grandes études de Paganini, S.141: No. 6 |
| 62. Lyric Pieces Op.47 Grieg | 92. S. 1413 in G \sharp Minor, La Campanella |
| 63. Lyric Pieces Op. 54 No. 4 | 93. S.541 No.3 in A \flat Major |
| 64. Morning Mood from Peer Gynt Suite No. 1 | 94. Symphony No.10 - I.
Adagio Complete Score |
| 65. Sonata in E Minor, Hob. XVI: 34 (I: Presto) | 95. Song Without Words Op.85 No.3 |
| 66. String quartet - Op.76, No.5, in D major | 96. Song without Words Op. 38 No.6 |
| 67. Cello Concerto C Major Movement 1 | 97. Song Without Words Op.30 No.5 |
| 68. Piano Sonata in F Major HOB.XVI/23 | 98. Melodie Op.4 No.2 in C minor |
| 69. Haydn Sonata Hob. XVI37 Mov. 1 D Major | 99. Songs without words Op.30 No.1 |
| 70. String Quartet Op.64 No.3 | 100. Songs Without Words Op.19 No.6 |
| 71. Piano Concerto in D major | 101. Songs Without Words Op.62 |
| 72. Die Schöpfung Mit Würd' und Hoheit ange-
tan | 102. Wedding March |
| 73. Piano Sonata in E minor HOB. XVI/34 | 103. Songs Without Words Op.19 No.4 |
| 74. Sonata in C minor HOB/XVI:20 | 104. Song Without Words Op.19b No.1 |
| 75. String Quartet in C major ("Emperor") Op.
76 No. 3 | 105. Songs Without Words Op.19 No.3 |
| 76. Die Fledermaus Grunfeld Op. 56 Konzert-
paraphrase | 106. Piano Sonata No.1 K.279 |
| 77. Radetzky March | 107. Sonata No. 5 1st Movement K.283 |
| 78. Pizzicato Polka Arranged for Piano Solo | 108. Sonata No. 7 1st Movement K. 309 |
| 79. The Blue Danube Accordion Solo | 109. Piano Sonata No.8 in A minor K.310300d |
| 80. Tratsch-Polka Op.214 | 110. Piano Sonata No. 8 in D Major,
K. 311 (284c): I. Allegro con spirito |
| 81. Strauss Die Fledermaus Op.362 Overture | 111. Piano Sonata No.18 in D major K |
| 82. Hungarian Rhapsody No. 2 | 112. Piano Concerto No.23 in A major K.488 |
| 83. Etude S.136 No.4 | 113. Mozart Rondo Alla Turca |
| 84. Trois Etudes de Concert No. 3 | 114. Piano Sonata No. 16 - Allegro |
| 85. Der Müller Und Der Bach. D795, S.5652 | 115. Sonata No.11 in A major K.331 |
| 86. Hungarian Rhapsody No. 6 | 116. Pictures at an Exhibition: No.2, Il
vecchio castello |
| 87. Etude S.136 No.5 | 117. Pictures at an Exhibition 13 8. Catacombae |
| 88. Etude S.136 No.9 | 118. Pictures at an Exhibition 14 Cum mortuis in
lingua mortua |
| 89. William Tell Overture Finale | 119. Pictures at an Exhibition Movement 15
(No.9) |
| 90. Romance S.169 | |

120. Pictures at an Exhibition 16 10
121. Pictures at an Exhibition-Gnomus (The Gnome) & Promenade
122. Pictures at an Exhibition
123. Strauss Die Fledermaus Op.410 Overture
124. Piano Concerto in G major - II
125. Gaspard de la Nuit, No.2, "Le Gibet"
126. Gaspard de la Nuit, No. 1, "Ondine"
127. Flight of the Bumblebee Piano
128. Concerto No.1 in A minor
129. Sans The Cuckoo in the Depths of the Woods
130. Sans - Fossils Transcribed for Piano
131. 2nd Piano Concerto 1st Movement Piano solo
132. Introduction and Rondo Capriccioso Op.28
133. Le Cygne The Swan Easy Piano by Free MusicKey
134. Allegro Appassionato Cello Piano
135. Piano Sonata D.784 - 1st movement
136. Impromptu in C minor No.1 Op.90
137. Impromptu No. 3 Op. 90 D 899 G \flat Major Transcription
138. Impromptu No.3 Op.90 D 899 G Major Transcription de Liszt
139. Piano Sonata No.19 in C minor
140. Sonata Op.42
141. Ave Maria
142. Winterreise D.911 No.1 - Gute Nacht
143. Winterreise D.911 No.5 - Der Lindenbaum
144. Die Forelle D. 550 Op. 32
145. Winterreise D.911 No.24 - Der Leiermann
146. Schwanengesang D.957 No.4
147. Waltz Op. 18 no. 6 in B minor
148. Piano Sonata No.2 in G
149. Vers La Flamme Op.72
150. Etude Opus 8 No. 12 in D Minor

B Appendix: Sample Questions

Using *Solfeggetto in C minor* by C. P. E. Bach as an illustrative example, the following section presents sample questions from each of the four levels in MSU-Bench. In total, the dataset comprises 1,800 QA pairs drawn from the 150 complete musical scores, with 450 questions allocated to each level. The structure of these questions is exemplified below. Here, **Q** denotes the question and **A** denotes the reference answer.

Level 1

- **Q1:** Is the piece written with an anacrusis (upbeat bar)?
A1: No.
- **Q2:** What is the tempo in beats per minute?
A2: 150 bpm.
- **Q3:** What is the initial key?
A3: C minor.

Level 2

- **Q1:** In bar 1, what is the dynamic marking?
A1: f.
- **Q2:** In bar 11, is there an accidental, and what is it?
A2: F \sharp 4, A \natural 4.
- **Q3:** In bar 25, what is the ornament on note D2?
A3: Passing note, neighbour note.

Level 3

- **Q1:** In bar 33, what is the chord progression? (use I, II, III, IV, V, VI, VII)
A1: I–V7.
- **Q2:** In bar 27, what scale degree is the first chord? (use I, II, III, IV, V, VI, VII)
A2: I.
- **Q3:** In bar 12, is there a dominant chord?
A3: No.

Level 4

- **Q1:** What is the predominant rhythm of the piece?
A1: Semiquaver.
- **Q2:** In bar 1, in which register is the melody?
A2: Middle register.
- **Q3:** What are the main features of the motif?
A3: Third interval.

C Appendix: Instructions Given to Annotators

As the annotators are experienced, we provide instructions only on how to format reference answers, as follows:

1. For enquiries regarding ornaments, identify the category of the ornament, such as “appoggiatura” or “mordent”, rather than listing the specific auxiliary pitches;
2. Utilise Scientific Pitch Notation (SPN) for all pitch identifications. Ensure that the octave number is included immediately after the note name (such as C4 rather than C);
3. For harmonic analysis, represent chords using Roman numerals to denote scale degrees and qualities (such as I, V7, or iv) rather than using macro-analytical symbols or jazz lead sheet notation;
4. When identifying keys, format the response with the capitalised note name followed by the mode in lowercase, formatted as “[Note] major” or “[Note] minor”;
5. If a single question elicits multiple distinct answers, separate each item with a comma to facilitate accurate verification.

D Appendix: Recruitment and Payment

We recruited ten doctoral candidates specialising in music from leading conservatoires in China. All participants began their formal musical training in early childhood, and consequently, each individual possessed approximately two decades of musical experience. Participation in this study was entirely voluntary and uncompensated. The decision to forgo monetary compensation was deemed adequate given the status of the participants as junior experts intrinsically motivated to contribute to academic discourse, rather than crowdsourced workers requiring financial incentives. This volunteer-based model is consistent with the norms of academic participation for this specific demographic in the region and adheres strictly to the ethical guidelines approved by the host institution.

E Appendix: Detailed Breakdown of Question Types for Levels 1–4

In this section, we provide a comprehensive breakdown of the question types encompassed within

each level of MSU-Bench, as illustrated in Figure 4.

E.1 Level 1

Onset Information. The questions ask about onset information that appears at the very beginning of a musical score, including the composer and title, the initial key, clef, and time signature, the presence of an anacrusis, the instruments involved, as well as the opening tempo, metronome marking, and expression indications. Collectively, these elements establish the basic identity, notation, and performance instructions that frame how the piece is read and interpreted from the outset.

E.2 Level 2

Notation & Note. Level 2 questions differ from Level 1 questions in both scope and depth. Whereas Level 1 focuses on onset information that is immediately visible at the start of a score, such as composer, title, initial key, time signature, clef, instrumentation, tempo, and expression indications, Level 2 shifts attention to localised details within the body of the music. These questions aim to examine specific bars for note-level features (highest or lowest pitch, presence of accidentals, shortest or longest note and rest values), performance instructions (dynamics, articulation, ornaments, tempo changes, expression markings), and structural signs (repeat signs, clef or key changes, modulation).

E.3 Level 3

Harmony & Chord. Level 3 questions delve into mid-level musical structures, focusing on harmonic and chordal analysis within specific bars. Level 3 progresses from the recognition of individual symbols to the analytical reasoning required for understanding harmonic structure. Rather than focusing on isolated surface features, such as a dynamic marking or an accidental within a single bar, these questions require the interpretation of tonal function. LLMs or VLMs are expected to identify chords by scale degree using Roman numerals, determine inversions and spacing, recognise cadences including perfect, imperfect, and interrupted, and distinguish non-chord tones such as suspensions (sus.) and anticipations (ant.). In addition, this level addresses harmonic progressions extending across multiple bars, examines whether passages commence or conclude on the tonic, and considers

the treatment of chord tones that are omitted or repeated.

E.4 Level 4

Texture & Form. Level 4 extends beyond the recognition of notation and harmony to encompass piece-specific understanding and large-scale structural analysis. The questions require the identification of a work’s genre and performance medium, including its orchestration or ensemble, as well as its formal design, such as principal and secondary themes, transitions, and main sections. They also address thematic materials by asking where a motif first appears, its defining features in terms of rhythm, register, or instrumental part, and the ways in which it is developed or ornamented. Further areas of focus include prevailing textural and accompanimental conventions, the predominant tempo and rhythmic character of the work, and, on occasion, the number of movements. In the case of instrument-specific repertoire, idiomatic features such as bowings are considered. Overall, Level 4 demands the synthesis of information across extended spans of music in order to characterise style, form, texture, and thematic organisation, thereby moving towards holistic musical analysis rather than bar-by-bar observation.

F Appendix: Distribution of Composers, Periods, and Genres in MSU-Bench

We provide a detailed account of the composer coverage, historical periods, and genre distribution of the musical scores included in MSU-Bench. The purpose of this analysis is to show that the benchmark is not narrowly tailored to a single stylistic tradition, but instead reflects a broad and representative spectrum of Western art music commonly encountered in formal music education and musical analysis.

F.1 Composer Distribution

MSU-Bench encompasses works by a diverse set of composers who occupy central positions in the Western canon, as shown in Figure 7. The selected scores include representative works by Johann Sebastian Bach, Ludwig van Beethoven, Frédéric Chopin, Johannes Brahms, Claude Debussy, Modest Mussorgsky, Hector Berlioz, Robert Schumann, Franz Liszt, and Pyotr Ilyich Tchaikovsky, among others. This selection strategy ensures coverage of composers associated with distinct compositional

languages, formal conventions, and notational practices. As a result, the benchmark exposes models to a wide range of stylistic fingerprints, from contrapuntal writing and functional harmony to extended chromaticism and colouristic textures.

F.2 Historical Periods

As shown in Figure 8a, the musical scores in MSU-Bench span a broad historical range, covering the Baroque, Classical, Romantic, and twentieth-century repertoire. The distribution is weighted toward Romantic music, while still maintaining meaningful representation from earlier and later traditions. This temporal spread exposes models to contrasting compositional idioms, including contrapuntal writing, formal balance, chromatic expansion, and more fluid twentieth-century harmonic and textural practices. Such diversity is important for evaluating whether models can generalise musical reasoning beyond a single historical style.

F.3 Genre and Formal Categories

As shown in Figure 8b, MSU-Bench covers a broad range of genres and formal categories, including sonatas, character pieces, fugues, waltzes, nocturnes, etudes, rhapsodies, symphonies, concertos, art songs, and others. These genres differ not only in surface style but also in their structural and textural demands: contrapuntal works require attention to voice leading and imitation, whereas sonatas, symphonic movements, and concertos more often involve large-scale formal reasoning. This distribution reflects the coexistence of extended forms and shorter self-contained works, allowing the benchmark to assess not only local notational understanding but also broader reasoning about form, thematic organisation, and texture across diverse compositional contexts.

G Appendix: Inference and Evaluation

Inference and evaluation are conducted entirely on a local GPU environment (A800). We use vLLM (Kwon et al., 2023) as the inference backend and adopt the unified data loading and prompt templating mechanisms provided by LLaMA-Factory (Zheng et al., 2024) to ensure a consistent inference pipeline. The evaluated models include both unmodified base models and parameter-efficiently fine-tuned variants using LoRA. Whether a LoRA adapter is loaded during inference is explicitly controlled by configuration settings, ensuring fair comparability across different model variants.




Figure 7: Frequency distribution of composers represented in MSU-Bench. The histogram illustrates the number of pieces per composer, with Franz Schubert, Liszt, Mendelssohn, and Edvard Grieg appearing most frequently, while representation gradually decreases for others.

During inference, the model generates answers autoregressively given the musical score images and the corresponding question text. Random sampling is used for decoding, including temperature sampling as well as top-p and top-k sampling. In our experiments, the decoding parameters are consistently set to temperature = 0.95, top-p = 0.7, and top-k = 50.

During the LLM-as-a-judge evaluation, for each sample, the judgement models take a predicted answer and a reference answer as input and produce a binary decision (1 for correct, 0 for incorrect). To ensure determinism, a fixed prompt template in Appendix H is used, and the decoding temperature is set to 0, allowing only a single judgement token to be generated.

Across all evaluated samples, the inter-judge agreement among ChatGPT-5, Claude Sonnet 4, Gemini 2.5 Pro, and an independent human assessment is 83.33%, with a 95% bootstrap confidence interval of (79.44%, 87.22%). This high unanimity rate indicates substantial consistency across judges and supports the reliability of the automatic evaluation pipeline.

H Appendix: Prompt Details

This appendix presents the exact prompt templates used throughout our inference and evaluation pipeline. The prompts are designed to support both general music-theoretic questions and score-localisation-based questions, while also providing a standardised instruction format for the LLM-as-a-judge setting. The same template is applied across models to ensure consistency in prompting and evaluation. For localisation-based questions, the

model is instructed to identify the relevant musical position before producing an answer, thereby reducing unsupported guessing and mitigating hallucinations in position-sensitive tasks. For the judgement setting, the prompt enforces a strict binary decision based on semantic equivalence between the predicted answer and the reference answer.

H.1 Inference User Prompt

```
Here is a piece of music in {score_modality}:
{score}

Question:
{Question}
```

H.2 Inference System Prompt

```
You are a professional assistant specialised in
music theory and musical scores.
You can answer both conceptual music questions
and score-based questions.
If a question requires locating a specific
position in a musical score
(such as a chapter, movement, section, bar,
measure, beat, staff, or voice),
you must first identify that position in the
score and then read the required
information from it.
```

Rules:

- Determine whether the question is score-localisation-based or not.
- For score-localisation questions, do NOT guess or generalise.
- For non-localisation questions, answer using standard music theory knowledge.
- If the required information cannot be determined with certainty, output "Unknown".
- Do NOT provide unnecessary explanations unless explicitly asked.

```
Your output should be accurate, concise, and
directly answer the question.
```




Figure 8: Distribution of musical periods and genres in MSU-Bench. (a) shows the historical periods of the selected scores, ranging from Baroque (5%) to Impressionism (7%) and various stages of Romanticism, with Late Romantic (19%) and Early Romantic (18%) being most prominent. (b) presents the genre distribution, where large-scale works (31%) and expressive pieces (25%) constitute the majority, followed by contrapuntal (11%), dance-based (8%), and other categories.

H.3 LLM-as-a-Judge Prompt

You are an expert in music theory. Your task is to evaluate the correctness of a Predicted Answer based on a Reference Answer.

Reference Answer: "{reference}" Predicted Answer: "{prediction}"

Instruction:

Relying on your music theory knowledge, compare the semantic meaning of the Predicted Answer to the Reference Answer.

If the Predicted Answer matches the Reference Answer, output 1.

If the Predicted Answer contradicts or fails to match the Reference Answer, output 0.

Do not output any text or explanation. Output only the single digit.

Decision:

I Appendix: Different Training Settings

To investigate the effect of different input modalities on model adaptation, we design three distinct fine-tuning strategies for Qwen2.5-VL-3B-Instruct. **PDF.** In this setting, we treat PDF sheet music as the visual input modality. The model receives images rendered from PDF pages, and both the visual

encoder and the language model are updated during training. This setting evaluates the model’s ability to extract structural and symbolic information directly from visual sheet-music representations.

ABC notation. Here, we replace PDF images with ABC notation as the only training input. Since this modality does not require visual parsing, we freeze the visual encoder to reduce computational overhead and update only the language model and LoRA adapters. This strategy evaluates whether ABC notation alone is sufficient for enabling VLMs to learn music-theoretical patterns.

PDF + ABC notation. In the multimodal setting, we provide both PDF images and ABC notation for each score. Both the visual encoder and the language model are fine-tuned. The objective is to examine whether complementary information from visual and symbolic modalities produces better performance than unimodal training. By integrating structural cues from PDF files with explicit symbolic tokens from ABC, this approach is expected to enhance robustness and generalisation across diverse tasks. However, the combination of both modalities constrains the maximum number of tokens available for training. The affected long scores are removed from the training pool only before the train/validation split is finalised; the fixed 30-score test set is never used for training and is unaffected




Figure 9: The inference time for models exceeding 40% overall accuracy.

Table 5: More zero-shot evaluation results on MSU-Bench (full), with the highest accuracy in **bold**. We evaluate 12 questions per score in a single run for each model to report the accuracy for each level and overall.

Models	Musical Score Understanding Benchmark				
	Level 1 (450)	Level 2 (450)	Level 3 (450)	Level 4 (450)	Overall (1800)
<i>Textual QA</i>					
Qwen3-4B	20.00	10.00	8.67	13.11	12.94
Qwen2.5-VL-7B-Instruct	32.00	9.11	17.56	13.33	18.00
Qwen2.5-VL-72B-Instruct	34.44	18.00	18.89	12.67	21.00
Llama 4 Scout	48.44	25.78	26.89	26.44	31.89
Qwen2.5-VL-32B-Instruct	50.67	20.22	26.22	37.56	33.67
Gemini 2.5 Flash	50.22	31.11	30.67	24.89	34.22
Qwen3-Next-80B-A3B-Instruct	57.11	23.11	25.33	34.00	34.89
Deepseek-V3	52.89	32.67	30.22	29.56	36.33
Llama 4 Maverick	52.67	31.56	28.44	33.56	36.56
Qwen3-Max	54.67	31.56	31.78	40.67	39.67

by these exclusions. The excluded training scores are:

1. *Piano Sonata No. 5 in C Minor, Op. 10 No. 1*
2. *Sonate pour Violoncelle et Piano*
3. *Piano Concerto in A Minor, Op. 16*
4. *Hungarian Rhapsody No. 2*
5. *Ballade No. 1 in G Minor, Op. 23*
6. *Ballade No. 4 in F Minor, Op. 52*
7. *Sonata, Op. 42*
8. *Concerto No. 1 in A Minor*

J Appendix: More results for Baselines

In Table 5, we report additional zero-shot baseline results on the full MSU-Bench benchmark to provide a broader comparison beyond the

strongest models discussed in the main text. Models such as Qwen3-Max and Llama 4 Maverick remain below 40% overall accuracy, indicating that strong general-purpose language capabilities do not directly translate into robust musical score understanding. Among the open-source models evaluated, Qwen3-VL-235B-A22B-Instruct demonstrates the strongest overall performance, exceeding the text-only Qwen3-Max by about 4%. This result suggests that large-scale multimodal modelling provides a measurable advantage even when performance is still far from saturation. In contrast, smaller models such as Qwen3-4B and Qwen2.5-VL-7B-Instruct perform considerably worse, thereby highlighting the limitations of lightweight architectures in zero-shot musico-logical reasoning tasks. The level-wise breakdown further shows that performance gains are not confined to a single difficulty level, but instead reflect more stable improvements across the benchmark as model capacity increases.



Figure 10: Representative score excerpt for the Level 2 and Level 4 PDF-QA failure cases.

K Appendix: Inference Time

We report the wall-clock inference time required by models achieving more than 40% overall accuracy in the zero-shot setting, corresponding to the results in Table 2 and Figure 9. The evaluation is conducted on the full MSU-Bench.

Overall, a clear trade-off between performance and computational cost is observed. Large proprietary models, including Gemini 2.5 Pro, ChatGPT-5, and ChatGPT-5-mini, attain the highest accuracies but require substantially longer evaluation times, typically exceeding eleven hours for a complete run. This latency arises from a combination of model size, inference overhead, and limited parallelism in sequential generative question answering.

In contrast, Qwen3-VL-235B-A22B-Instruct achieves a competitive overall accuracy of 41.22% while requiring only approximately one hour of evaluation time. This result highlights the efficiency advantages of certain open-source architectures, which deliver favourable accuracy–efficiency trade-offs under identical evaluation protocols.

These findings suggest that inference time constitutes a practical constraint that should be considered alongside raw accuracy. Efficient models may offer a more viable solution for rapid development and comparative analysis, even when absolute performance differences are relatively small.

L Appendix: Representative PDF-QA Failure Cases

To complement the aggregated visual-QA results, we present three representative PDF-QA failure cases on Johann Sebastian Bach’s Fugue in D minor, BWV 948, generated by Claude Opus 4 (Figures 10 and 11). These examples highlight common error patterns in full-score reasoning, particularly incorrect bar-level grounding and hallucinated local details. In the following examples, **Q** denotes the question, **P** the model prediction, and **GT** the ground truth.

L.1 Level 2 failure case

Q: In bar 15, what is the highest pitch?

P: Looking at bar 15 on page 2, the right hand (treble clef) begins with some notes featuring sharps and then moves through a melodic line. The highest notes in the treble staff in bar 15 appear to reach up to approximately D5 (the D above middle C, sitting just below the top of the treble staff). There are notes on the upper portion of the staff, but nothing ventures above the staff lines into ledger-line territory. The highest pitch in bar 15 is D5.

GT: A4.

Observation: The model misreads the local pitch content in the target bar and overestimates the upper note range, indicating weak note-level grounding even when the queried bar is explicitly specified and the relevant notes are visually accessible.



Figure 11: Representative score excerpt for the Level 3 PDF-QA failure case.

L.2 Level 3 failure case

Q: In bar 54, is there a perfect cadence?

P: It doesn't render the actual notes clearly enough for me to read the precise pitches and harmonies. What I can see from the layout is that bar 54 sits in the middle of a passage on page 6, but I can't reliably determine the specific chords to confirm whether there's a V-I (perfect/authentic) cadence at that point.

GT: No.

Observation: Instead of resolving the harmonic relation in the target bar, the model fails to ground the local notation strongly enough to make a cadence judgement and abstains from answering.

L.3 Level 4 failure case

Q: In which bar does the melodic motif first appear?

P: The fugue subject (the main melodic motif) first appears in bar 2 in the right hand — it's the opening solo entry that begins with a rest on beat 1, then launches into the characteristic rising figure followed by running sixteenth notes. This is the subject that gets developed throughout the fugue. Bar 1 is essentially empty (a full-bar rest), so the motif's first appearance is in bar 2. Subsequent entries of the subject appear in various bars as the fugue unfolds (e.g., the bass voice enters later around bar 11–12 with the answer).

GT: Bar 1.

Observation: This error reflects incorrect bar anchoring at the opening of the piece: the model imposes an inferred rest-based reading and shifts the first motif entry from bar 1 to bar 2.

M Appendix: One vs Batched Questions

In Figure 12, we compare two evaluation settings for VLMs on MSU-Bench: asking each question individually and asking all 12 questions associated with a score in a single batch. The results reveal substantial variation across models, with batched prompting generally leading to higher overall accuracy. This pattern suggests that jointly presenting the questions allows models to better exploit shared score-level context, especially when lower-level observations can support answers to higher-level questions. At the same time, the magnitude of the gain differs across models, indicating that some systems benefit more than others from cross-question contextualisation within a single inference run.

N Appendix: Analysis of LSR for Models Adapted Using LoRA

We also report LSR for models adapted with LoRA, complementing the zero-shot LSR analysis in Figure 6b. LSR provides a stricter view of model behaviour than per-question accuracy, because a




Figure 12: Overall accuracy comparison under the Question-by-Question and Single-Run (Batch) evaluation for VLMs. The horizontal axis reports overall accuracy in percentage, while each bar corresponds to a specific model.

score is counted as successful at Level l only if all questions from Level 1 through Level l are answered correctly. As a result, it directly captures whether a model can maintain consistency across progressively more demanding stages of musical understanding rather than solving isolated questions in isolation.

Consistent with the zero-shot results, LSR decreases monotonically as the level increases, reflecting the growing difficulty of sustaining correctness across successive levels of musical comprehension. Even when a model performs reasonably well on individual questions, small errors at earlier levels accumulate and sharply reduce the number of scores that remain fully correct at later levels. This makes LSR especially informative for evaluating end-to-end score understanding.

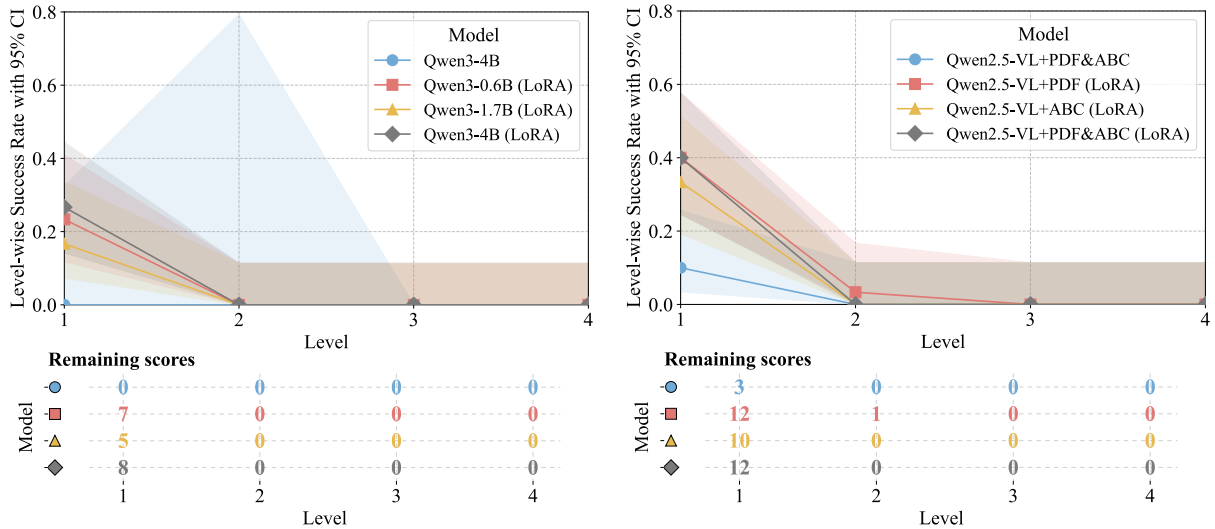
However, LoRA-adapted models exhibit markedly improved retention compared to their base counterparts, as shown in Figures 13a and 13b. For LLMs, adapted models maintain non-negligible success rates compared to the base model (Qwen3-4B), indicating that lightweight supervised adaptation helps preserve correctness across multiple levels rather than only boosting

isolated responses. Improvements are even more pronounced in VLMs. While zero-shot visual models fail almost entirely after Level 1, LoRA adaptation enables models such as Qwen2.5-VL-3B-Instruct to preserve a meaningful number of answerable scores at Level 1, especially when PDF and ABC inputs are combined.

Taken together, these results indicate that LoRA adaptation not only improves per-question accuracy but also enhances cross-level consistency, mitigating the compounding error effects revealed by LSR. Nevertheless, the sharp decline at higher levels persists across all settings, underscoring that robust, end-to-end musical reasoning over complete scores remains an open challenge even after adaptation.

O Appendix: MMLU

We evaluate the models adapted with LoRA and those without adaptation on MMLU in order to assess whether task-specific adaptation to MSU-Bench leads to catastrophic forgetting in general-domain knowledge. MMLU covers 57 subjects spanning Science, Technology, Engineering, and Mathematics (STEM), as well as the humanities, social sciences, and other professional domains,



(a) LLMs. Results for models trained solely on ABC notation without other modalities. The LoRA-finetuned models demonstrate a discernible advantage over the base Qwen3-4B model.

(b) VLMs. We compare VLMs trained on PDF, ABC, or both modalities against the baseline Qwen2.5-VL-3B-Instruct, which utilises combined inputs.

Figure 13: Level-wise Success Rate for Models Adapted Using LoRA.

Table 6: Evaluation of models conducted before and after LoRA on MMLU. Qwen2.5-VL-3B-Instruct is adapted using LoRA across the three input modalities described in Section 5.1.

Models	STEM	Humanities	Social Sciences	Other Subjects
Qwen3-4B	72.63	81.44	63.21	74.61
w/ LoRA	74.09(+01.46)	81.54(+00.10)	63.51(+00.30)	75.11(+00.50)
Qwen2.5-VL-3B-Instruct	60.60	75.63	58.72	69.65
w/ PDF	60.90(+0.30)	75.66(+00.03)	58.45(-00.27)	69.80(+00.15)
w/ ABC	60.47(-00.13)	75.79(+00.16)	58.13(-00.59)	69.62(-00.03)
w/ PDF&ABC	60.50(-00.10)	75.85(+00.22)	58.28(-00.44)	69.65(-00.00)

and therefore provides a broad reference benchmark for retained reasoning ability.

The results indicate that LoRA adaptation does not cause substantial degradation on MMLU. For the Qwen3-4B model, the application of LoRA yields consistent improvements across all reported categories, including STEM, Humanities, Social Sciences, and Other Subjects. This suggests that adaptation to the musical benchmark can in some cases improve general instruction-following or transfer behaviour rather than weakening it. In contrast, the performance of Qwen2.5-VL-3B-Instruct remains largely stable following adaptation with different input modalities, including PDF and ABC notation. The observed deviations are small, indicating that modality-specific adaptation does not materially erode the model’s general knowledge.

Overall, these findings suggest that the proposed fine-tuning setup is relatively robust with respect to forgetting: models can be adapted to musical

score understanding while largely preserving their competence on a broad external benchmark. This is important in practice, as it indicates that gains on MSU-Bench are not simply obtained by sacrificing more general reasoning ability.

P Appendix: Use of LLMs

In preparing this manuscript, we employ LLMs solely as an auxiliary tool for academic writing. Their use is restricted to *linguistic refinement*, including polishing grammar, improving clarity and fluency, and adjusting the structure and formatting of text and tables. We do not rely on LLMs for generating research ideas, designing methodologies, conducting experiments, performing data analysis, or interpreting results. All conceptual contributions, experimental designs, and substantive findings reported in this work are entirely our own. All final decisions regarding content, wording, and presentation are made by the authors.