



PianoCoRe: Combined and Refined Piano MIDI Dataset

DATASET ARTICLE

ILYA BOROVIK 

*Author affiliations can be found in the back matter of this article



ABSTRACT

Symbolic music datasets with matched scores and performances are essential for many music information retrieval (MIR) tasks. Yet, existing resources often cover a narrow range of composers, lack performance variety, omit note-level alignments, or use inconsistent naming formats. This work presents **PianoCoRe**, a large-scale piano MIDI dataset that unifies and refines major open-source piano corpora. The dataset contains 250,046 performances of 5,625 pieces written by 483 composers, totaling 21,763 h of performed music. PianoCoRe is released in tiered subsets to support different applications: from large-scale analysis and pre-training (**PianoCoRe-C** and deduplicated **PianoCoRe-B**) to expressive performance modeling with note-level score alignment (**PianoCoRe-A/A***). The note-aligned subset, **PianoCoRe-A**, provides the largest open-source collection of 157,207 performances aligned to 1,591 scores to date. In addition to the dataset, the contributions are: (1) a MIDI quality classifier for detecting corrupted and score-like transcriptions and (2) RAScoP, an alignment refinement pipeline that cleans temporal alignment errors and interpolates missing notes. The analysis shows that the refinement reduces temporal noise and eliminates tempo outliers. Moreover, an expressive performance rendering model trained on PianoCoRe demonstrates improved robustness to unseen pieces compared to models trained on raw or smaller datasets. PianoCoRe provides a ready-to-use foundation for the next generation of expressive piano performance research.

CORRESPONDING AUTHOR:

Ilya Borovik

Skolkovo Institute of Science and Technology, Moscow, Russia

ilya.borovik@skoltech.ru

KEYWORDS:

symbolic music, MIDI dataset, piano performance, musical score, score-performance alignment, expressive performance rendering

TO CITE THIS ARTICLE:

Borovik, I. (2026). PianoCoRe: Combined and Refined Piano MIDI Dataset. *Transactions of the International Society for Music Information Retrieval*, 9(1), 144-163.
DOI: <https://doi.org/10.5334/tismir.333>

1 INTRODUCTION

Musical scores and live performances are fundamental data sources for a wide range of music information retrieval (MIR) tasks. A score provides a symbolic representation of the written composition, while a performance captures a musician's unique interpretation through variations in timing, dynamics, and articulation. Modeling the relationship between these two domains is essential for analyzing the decisions performers make to convey musical structure and emotion to an audience. Furthermore, paired score-performance data enables computational expressive performance rendering, where trained models simulate human interpretation. For all these tasks, the scale, quality, and structure of available datasets are essential.

For piano music, numerous symbolic corpora have been developed to support computational performance analysis and modeling (Cancino-Chacón et al., 2018; Emerson and Harrison, 2025; Lerch et al., 2020). These resources fall into two categories. The first comprises high-fidelity recordings captured from computer-monitored acoustic pianos (e.g., Yamaha Disklavier) (Foscarin et al., 2020; Goebel, 1999; Hashida et al., 2018; Hawthorne et al., 2019; Hu and Widmer, 2023). The second category relies on automatic music transcription (AMT) (Benetos et al., 2018) to generate large-scale datasets from audio recordings (Bradshaw and Colton, 2025; Edwards et al., 2023; Kong et al., 2022; Lee et al., 2025; Zhang et al., 2022). While recorded datasets offer unparalleled expressive detail, they are often limited in scale and stylistic diversity. Conversely, AMT-based datasets provide diversity but often contain transcription errors and lack precise note-level alignments. Furthermore, incompatible naming schemes and metadata standards make it difficult to combine datasets without risking information leakage. Together, these challenges highlight a critical gap: a lack of a unified resource that combines the scale of transcribed data with the precision of recorded performances, all aligned to scores.

This gap is addressed by **PianoCoRe**,¹ a comprehensive dataset that combines and refines the largest open-source piano corpora of scores and performances. PianoCoRe contains 21,763 hours of piano music across 250,046 performances of 5,625 pieces by 483 composers, with scores available for 75.3% of performances. To make this data usable across diverse applications, it is released in tiered subsets:

- **PianoCoRe-C**: a complete mixed-source piano performance collection;
- **PianoCoRe-B**: a deduplicated and quality-assessed subset for large-scale pre-training;
- **PianoCoRe-A**: a subset containing performances note-aligned to scores; and
- **PianoCoRe-A***: a high-quality subset of the best-quality performances and note-level alignments.

Unlike previous efforts, PianoCoRe focuses on legal sustainability by restricting content to works in the public domain in the European Union, ensuring it remains a stable and sound resource for the academic community. To support diverse use cases, the dataset is archived on Zenodo² and mirrored on Hugging Face.³

By providing an annotated dataset that is larger and cleaner than previous resources, this work lays a foundation for the development of more intelligent computational piano performance models.

The main contributions of the work are:

1. The matching and combination of existing piano MIDI corpora into a single, large-scale unified collection with verified metadata (Section 3);
2. A deduplication and alignment-based heuristic for MIDI quality labeling and a trained classifier for filtering corrupted and inexpressive transcriptions, enabling the creation of the curated PianoCoRe-B dataset (Section 4);
3. RAScoP (Refined Alignment for Scores and Performances), a note-alignment refinement pipeline that cleans timing outliers, interpolates missing notes, and synchronizes performances with scores that has been used to produce the note-aligned PianoCoRe-A/A* subset (Section 5); and
4. An application of PianoCoRe to the task of performance rendering and a discussion of the benefits of the combined dataset for training compared to individual source datasets (Section 6).

The rest of this work is structured as follows: Section 2 reviews relevant piano datasets. Section 3 details the curation process for PianoCoRe. Section 4 introduces the MIDI quality classifier and deduplicated subset. Section 5 presents RAScoP and the note-aligned subsets. Section 6 evaluates PianoCoRe on expressive performance rendering. Finally, Sections 7 and 8 discuss limitations and conclude the work.

2 RELATED WORK

This section provides an overview of the most prominent piano score and performance datasets, categorized by primary data source and intended application. Table 1 provides a summary of the datasets relevant to PianoCoRe and statistics of PianoCoRe itself.

2.1 RECORDED MIDI PERFORMANCE DATASETS

One category of datasets consists of MIDI files captured directly from human performances on computer-monitored pianos (e.g., Yamaha Disklavier). These performances offer the highest fidelity of expressive detail at the symbolic level.

Dataset	Composers	Pieces	Performances	Hours	Sources	Scores	Alignments	Metadata
MAESTRO	43	–	1,276	199	R	no	no	P
(n)ASAP	16	222	1,067	92	R	100%	beat/note	P
GiantMIDI	2,786	10,855	10,855	1,237	T	no	no	S
ATEPP	25	1,596	11,742	1,009	T	43.6%	no	P, Q [†]
Aria-MIDI	19,021 [‡]	–	1,186,253	100,629	T	no	no	S, P [†]
PERiScoPe	82	2,738	46,473	3,784	R, T	81.9%	note	P [†]
PianoCoRe-C	483	5,625	250,046	21,763	R, T	75.3%	no	P[†]
PianoCoRe-B	478	5,591	214,092	18,757	R, T	75.0%	no	P[†], D, Q
PianoCoRe-A	151	1,591	157,207	12,509	R, T	100%	note	P[†], D, Q
PianoCoRe-A*	137	1,517	130,275	10,330	R, T-HQ	100%	note	P[†], D, Q

Table 1 Comparison of major symbolic piano performance datasets and **PianoCoRe** dataset with its tiers. Sources: R—recorded (Disklavier/Hardware), T—transcribed (Audio-to-MIDI), T-HQ—transcribed labeled as high quality. Metadata: P—performer, S—piano solo probability, D—deduplication flag, Q—quality label. [†]Annotations are not available for all performances. [‡]Number of unique composer names computed from raw metadata, not manually verified.

The **MAESTRO** dataset (Hawthorne et al., 2019) is the most influential in this category, with over 200 h of virtuosic performances from the International Piano-e-Competition. The high-quality, time-aligned audio–MIDI pairs have made it the standard for transcription benchmarks. However, its size and diversity are modest by modern deep learning standards.

The **ASAP** dataset (Foscarin et al., 2020) extends MAESTRO by adding musical scores and beat annotations. The dataset contains nearly 92 h of 1,067 performances from MAESTRO aligned at the beat level to 222 unique scores. Its extension, **(n)ASAP** (Peter et al., 2023), adds note-level alignments, making it the largest open-source recorded MIDI dataset with score-to-performance note alignments.

Several smaller curated datasets offer exceptional detail for specialized analysis tasks. The **Batik-plays-Mozart** corpus (Hu and Widmer, 2023) provides note-for-note alignments between professional MIDI performances of Mozart sonatas and expert-annotated scores. **Vienna 4x22 Piano Corpus** (Goebel, 1999) captures four classical music excerpts performed by 22 pianists. **SMD** (Müller et al., 2011) provides perfectly synchronized audio and MIDI for 50 performances of 50 pieces by 11 composers. **MazurkaBL** (Kosta et al., 2018) provides score-aligned beats, loudness, and expressive markings for 2,000 recordings of Chopin’s mazurkas. **CrestMusePEDB** (Hashida et al., 2018) contains 411 note-aligned performances of 35 classical pieces by 12 pianists. While invaluable for detailed study, these datasets’ narrow scope limits their utility for training general-purpose performance models.

2.2 LARGE-SCALE TRANSCRIBED MIDI DATASETS

To avoid the time-consuming process of collecting MIDI data recorded on sensor-equipped pianos, researchers increasingly use AMT (Benetos et al., 2018) to generate large datasets from publicly available audio.

GiantMIDI-Piano (Kong et al., 2022) was an early large-scale piano transcription effort (Kong et al., 2021), providing 1,237 h of classical piano MIDI across 10,855 pieces. The audio was sourced from performances of IMSLP repertoire downloaded from YouTube, covering compositions from a wide range of musical periods. However, GiantMIDI-Piano does not provide any musical scores, and the metadata contains duplicates and inconsistencies (see Section 3.3.3).

The **ATEPP** dataset (Zhang et al., 2022) captures 11,674 performances by renowned pianists, totaling over 1,007 h of transcribed music. About half of performances have a paired score without any note-level alignment. ATEPP provides quality labels (‘high quality’, ‘low quality’, ‘corrupted’) for some of the performances. However, as analyzed in Section 4.2, there are unlabeled corrupted transcriptions.

Aria-MIDI (Bradshaw and Colton, 2025) greatly expands the data scale dimension, offering over 100,629 h of transcribed piano music. Data was crawled, classified as piano solo, and annotated using a large language model-guided pipeline. The size of Aria-MIDI makes it valuable for self-supervised learning. However, the dataset lacks symbolic scores and complete annotations of musical pieces.

Other notable efforts include the **SUPRA** dataset (Shi et al., 2019), which digitized an archive of 52 h of 478

piano roll performances. In the piano jazz domain, the **PIJAMA** dataset (Edwards et al., 2023) provides 223 h of high-quality transcriptions of 2,777 performances by 120 pianists.

2.3 MIXED-SOURCE PIANO DATASETS

Although the above datasets are valuable, they exist in isolation, each with different structures and metadata conventions. Mixing them directly for piano performance modeling introduces the risk of information leakage between the training and test splits.

GigaMIDI (Lee et al., 2025) contains over 1.4 million MIDI files from diverse single- and multi-instrument sources, including ASAP, ATEPP, GiantMIDI-Piano, Vienna 422, SMD, and Batik-plays-Mozart. A valuable contribution is the set of heuristics for categorizing inexpressive MIDI data. However, unnormalized piece titles in GigaMIDI complicate piece-based grouping and comparison of the data.

The **PERiSCOPE** dataset (Borovik et al., 2025) represents an effort to bridge the gap between recorded and transcription-based MIDI datasets. It contains over 35,000 note-aligned score–performance pairs, matching and combining (n)ASAP and ATEPP with 2,158 h of web-collected audio transcribed to MIDI.

The described single-source and multi-source datasets face several limitations that **PianoCoRe** aims to resolve. First, collections often lack a standardized, easy-to-navigate directory structure and verified metadata, making them difficult to combine and extend. Second, datasets may pose legal risks due to the inclusion of modern, copyrighted works. Finally, MIDI transcriptions may be duplicated, corrupted, or transcribe musical score audios that provide no information for performance analysis and modeling.

3 PianoCoRe DATASET

This section details the construction of **PianoCoRe**. It presents a methodology for processing musical scores; matching works across diverse datasets; preprocessing the source files to resolve inconsistencies; and integrating them into a unified, navigable collection. The final dataset is presented at the end of the section.

3.1 NOTATION AND DEFINITIONS

The core entities and relations used throughout the manuscript and in the data collection and processing pipelines are as follows:

- **Note**, n : a MIDI note described by its pitch p , onset o , duration d , and velocity v : $n = (p, o, d, v)$. Notes are indexed $i \in \{1, \dots, N\}$ after sorting MIDI by onset, pitch, and duration;

- **Musical score**, y : a sequence of N_s score MIDI notes (y_1, \dots, y_{N_s}) ;
- **Performance**, x : a sequence of N_p performance MIDI notes (x_1, \dots, x_{N_p}) ; and
- **Alignment**, A : a sequence of score and performance notes pairs $\{(y_i, x_j) : y_i \in y \cup \{\emptyset\}, x_j \in x \cup \{\emptyset\}\}$, where $a_{ij} = (y_i, \emptyset)$ indicates a missing performed note and $a_{ij} = (\emptyset, x_j)$ – an inserted performance note. The number of matched notes (pairs with $y_i \neq \emptyset \wedge x_j \neq \emptyset$) is denoted as N_m .

The following four primary ratios are used to evaluate the relationship between a score and a performance:

- **Note Ratio**, R_n : a ratio of the number of notes between performance and score sequences:

$$R_n = \frac{N_p}{N_s} \quad (1)$$

Given the same musical content, note ratio identifies structural discrepancies, such as omitted repeats ($R_n \ll 1$) or transcription noise ($R_n \gg 1$);

- **Alignment Recall**, R_o : a proportion of score notes matched to the performance:

$$R_o = \frac{N_m}{N_s} \leq 1 \quad (2)$$

Recall represents the ‘completeness’ of the performance relative to the score;

- **Alignment Precision**, P_o : a proportion of performed notes matched to the score:

$$P_o = \frac{N_m}{N_p} \leq 1 \quad (3)$$

High precision indicates a clean performance with few noisy notes or insertions;

- **Adjusted Alignment Ratio**, R'_o : a relaxed quality metric that takes the highest of Recall (when $N_p \geq N_s$) and Precision ($N_p < N_s$):

$$R'_o = \frac{N_m}{\min(N_s, N_p)} = \max(P_o, R_o) \leq 1 \quad (4)$$

It ensures that a performance is not penalized for missing notes (e.g., skipped repeats) as long as the played notes match the score, and is not penalized for extra notes (e.g., transcription noise) as long as all score notes are present.

Furthermore, the two common types of symbolic errors handled during preprocessing are:

- **Duplicate Notes**: two or more notes having the exact same pitch, onset time, and duration and
- **Overlapping Notes**: a condition where a note n_i of pitch p starts while a previous note n_{i-1} of the same pitch is still active ($o_i < o_{i-1} + d_{i-1}$).

3.2 DATA-MATCHING METHODOLOGY

The essential part of a score and performance music dataset is the correct matching of scores and performances. One approach is to use composition entity resolution (Kong et al., 2022; Zhang et al., 2022) that compares the titles and available metadata for score and performance files. However, the music content may not reflect the title if the file is mislabeled or has a unique naming format.

MIDI-to-MIDI matching is used to combine datasets. This allows one to directly compare notes in musical scores and performances. It also enables one to match performances to musical scores that are only available in MIDI format and have no MusicXML (Good, 2001) counterpart. Finally, it allows one to match performances to other performances to obtain more labeled data when no scores are available.

3.2.1 Score processing

Before matching, the MusicXML files were converted to MIDI format using the `partitura` library (Cancino-Chacón et al., 2022) with the following refinements:

- **Dynamics and Tempo:** the `<sound>` tags and `dynamics` attributes for notes are processed to embed performance direction markings for dynamics and tempo directly into the note velocities and tempo changes of the score MIDI file.
- **Ornaments:** trills and mordents are unrolled based on the invisible notes available in MusicXML (`<cue/>` or `print-object="no"`). The base visible ornament note is removed to avoid overlapping note events.
- **Grace Notes:** acciaccatura and appoggiatura notes are expanded based on the definitions. Acciaccatura notes appear as a sequence of 32nd notes before the beat. Appoggiatura notes steal the duration of the main note.
- **Repeats:** for scores with repeats, two versions are created: a *maximal* version with all repeats unfolded and a *minimal* version with each repeat played only once (suffix `_mini` in the file name).

These changes ensure fair consideration of score structure and performance-specific elements in MIDI score files. To simplify the management of the created dataset, the full set of possible repeat structures in the scores was not considered.

3.2.2 Candidate pair selection

To avoid a brute-force comparison of all files, a filtering step to identify a smaller set of candidate pairs is performed. A score is paired with a performance if they meet the following criteria:

- **Composer:** the composer names, extracted from file paths or metadata tags, match;

- **Note Count:** the note ratio R_n falls within a plausible range of close length: $0.75 \leq R_n \leq 1.33$; and
- **Keywords:** if available, the catalog numbers, and key/scale information within the titles match.

This pre-filtering enables efficient application of computationally intensive, alignment-based verification.

3.2.3 Note alignment and verification

For the final step, note-level alignments for candidate pairs were computed using the `DualDTWNoteMatcher` from Paragonar (Peter, 2023). The underlying dynamic time warping (DTW) implementation was optimized using Numba's just-in-time (JIT) compilation (Lam et al., 2015). The optimized version works, on average, 12 times faster on the ASAP dataset. This optimization was essential for performing millions of pairwise alignments within a reasonable timeframe.

A candidate pair is considered a definitive match if the alignment achieves $R_o > 0.7$ (more than 70% of score notes matched to the performance). This threshold was chosen empirically to ensure a global overlap between the sequences with close score and performed repeat structures. Unmatched notes may correspond to omitted repeats, transcription errors, or specific interpretations. These data are still valuable for performance-only applications, including large-scale pre-training.

Performances that fail to align with the maximal unfolded score are matched to the minimal one, increasing data retention. The exact repeat structure of the performances is not detected. For trills, the number of notes may differ between performances and scores. However, unrollment of trills in the score MIDI yields a higher alignment recall than aligning multiple performed notes to a single base trill note.

Alignments are stored in compressed `.npz` files compatible with the original MIDI files. Each file contains arrays describing the attributes of the aligned score and performance notes: indices, pitches, and onset/offset times. Insertions and deletions are represented by the sentinel value `-1` for missing attributes.

3.3 SOURCE PERFORMANCE DATASETS

PianoCoRe is built by refining and integrating open-source piano MIDI datasets. This section describes the steps taken to improve the quality of source datasets before combining them under a single collection.

3.3.1 ASAP dataset

The (n)ASAP dataset v2.1.1 (Peter et al., 2023)⁴ was used. The original score MIDI files, exported using MuseScore (Watson, 2018), contain data-parsing issues like unrealistic time signatures (e.g., 65/4, 25/32), cut measures with anacrusis, duplicated notes, and notes with zero duration. These were corrected by re-generating score MIDI files using the standardized pipeline (Section 3.2.1). The

performance MIDI files were cleaned by removing duplicate notes, truncating durations of the first of the two overlapping notes (such that $o_i = o_{i-1} + d_{i-1}$), and removing all notes shorter than 5 ms. There are 208 score and 94 performance MIDI files with zero duration notes in the original dataset.

3.3.2 ATEPP dataset

The ATEPP v1.2 dataset (Zhang et al., 2022)⁵ was used. Only 5,091 of 11,674 transcribed performances are paired with scores without an alignment. ATEPP shares the scores with ASAP, but not all suitable scores (e.g., the entirety of Chopin) are present in ATEPP. By matching two datasets, 39 scores from ASAP can be assigned to 827 performances in ATEPP.

As a preprocessing step, score MIDI files were computed from MusicXML files, similar to scores in ASAP. Also, the following metadata issues were corrected: merging duplicate movements under different names (49 movements and 265 reassigned performances), performances with a wrong piece name (24 movements and 43 performances), and performances without a score in the metadata (3 scores and 14 performances). These problems were fixed by matching and checking performances and scores of the same composer.

3.3.3 GiantMIDI-piano

For GiantMIDI-Piano (Kong et al., 2022), a curated subset of the original data⁶ consisting of 7,236 MIDI files was used. The analysis of the metadata showed duplicates (by YouTube ID) in the original curated data. In total, 315 MIDI transcriptions were distributed under multiple composition names. Also, manual inspection during the matching process revealed other inconsistencies. A MIDI file may represent only a specific movement of the annotated piece, or it may be a performance of a different piece mistakenly matched after a YouTube search.

Since checking and annotating all MIDI files is exhaustive, only sequences that matched with the scores and performances from other examined datasets were used. The final subset included 2,139 performance MIDI files of musical pieces by 402 composers.

3.3.4 PERiSCOpe

The PERiSCOpe v1.0 dataset (Borovik et al., 2025)⁷ was processed by excluding performances from ASAP or ATEPP. Only the remaining 34,773 performance MIDI files transcribed from audio sources using Transkun V2 (Yan and Duan, 2024) were used. The dataset required no specific process except for common transcription artifacts, described below in Section 3.3.6.

3.3.5 Aria-MIDI

From the Aria-MIDI v1 dataset (Bradshaw and Colton, 2025) with 1,186,253 transcribed MIDI files,⁸ 621,132

files that had a composer in the metadata were filtered and used. There are 19,021 unique composer names in the filtered subset.

An important difference in Aria-MIDI is how sustain pedals are encoded. The transcribed files do not distinguish between pressed and sustained note durations. The durations were predicted as sustained even when the sustain pedal was predicted separately.

3.3.6 Transcription artifacts

One issue fixed for all transcribed MIDI datasets is the error with ‘infinite’ pitches, where notes span until the end of the file. This artifact arises when open-source transcription models (Kong et al., 2021; Yan and Duan, 2024) produce unmatched note-on and note-off events due to offset or sustain-pedal decoding errors. During MIDI serialization, such notes remain active till the end of the sequence. An algorithm to identify and correct note durations was developed to repair performances in the source datasets: ATEPP (30 MIDI files), GiantMIDI (9), PERiSCOpe (92), and Aria-MIDI (5,501).

3.4 MUSICAL SCORE DATA SOURCES

To maximize the number of aligned performances, the score library was expanded beyond ASAP and ATEPP and included public domain MusicXML scores from the PDMX dataset (Long et al., 2025), originally sourced from MuseScore.⁹ In addition, the sequenced MIDI scores from KunstderFuge¹⁰ and ClassicalMIDI¹¹ websites were used solely for enriching the representation of annotated performed compositions in PianoCoRe. The copyrighted scores are not redistributed in the final dataset. Since KunstderFuge provides live performance and orchestral MIDI files, inexpressive solo piano sequences were filtered out using a Note Onset Median Metric Level (NOMML) heuristic from GigaMIDI (Lee et al., 2025). Finally, during the iterative data-matching process, 421 public domain scores from MuseScore were manually sourced for the most frequently performed compositions that lacked a score.

3.5 DATA-COMBINATION PROCESS

The **PianoCoRe** dataset was assembled using a semi-automated, iterative process designed to merge multiple sources into a single, structurally unified collection. This process relies on the data matching and note alignment (Section 3.2), supplemented by manual curation and labeling to resolve ambiguities.

The main strategy was to establish a unified data organization and gradually integrate scores and performances from source datasets. The combination process unfolds in three stages, illustrated in Figure 1:

- 1. Core Structure:** The process began with the merging of two foundational datasets: ASAP and ATEPP. Performances and scores from ASAP were matched

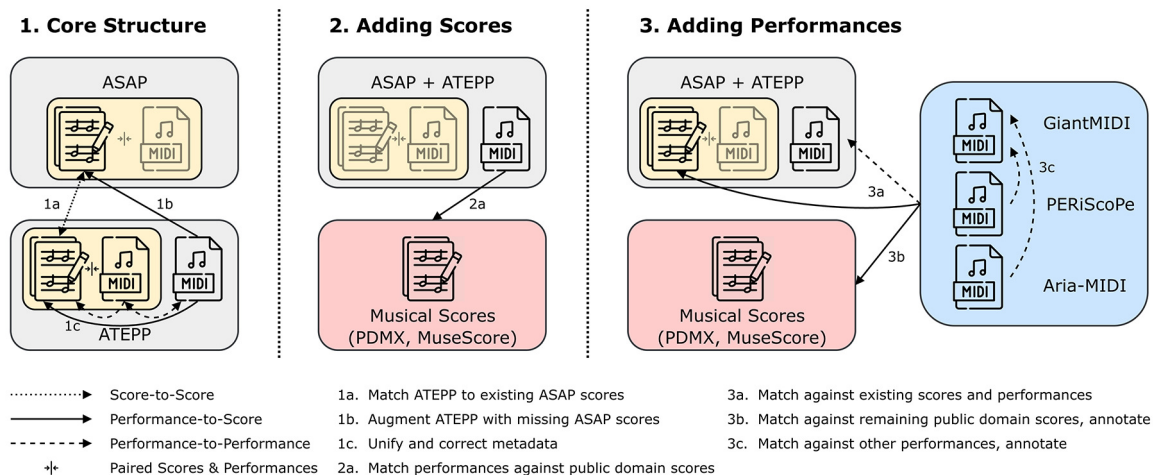


Figure 1 The three-stage data matching and annotation pipeline used to create PianoCoRe dataset.

and reorganized into the unified ATEPP directory structure. Lastly, the 21 ASAP pieces not present in ATEPP were distributed under new directories. This created a unified base of recorded and transcribed performances with their corresponding scores.

2. Adding Scores: The core dataset was then augmented by matching its performances against a large corpus of scores from PDMX, KunstderFuge (KDF), and Classical MIDI (CM), along with manually added MuseScore (MS) files.

3. Adding Performances: The final step involved the integration of the performance datasets: GiantMIDI-Piano, PERiScoPe, and Aria-MIDI. Performances were matched against available scores based on the initial candidate pair selection. If a piece was not present in the dataset, a new directory containing the score and matched performances was added. To further increase data coverage, remaining performances were matched against those without a score from ATEPP and against each other to identify additional composition-based links.

Throughout the process, automated matches were reviewed. For new pieces, composition and movement titles were manually verified and standardized using IMSLP¹² and web search. This step ensured consistency, corrected mislabeled files, and prevented compositions from being cataloged under different names. To ensure compliance with copyright standards, only works in the public domain in the European Union¹³ were included.

3.6 PianoCoRe-C DATASET

The result of the data combination is **PianoCoRe-C** dataset, where ‘C’ stands for ‘Core’ or ‘Combined’. This dataset represents the most diverse collection of piece-wise annotated piano performances. It contains 250,046 performance MIDI files for piano pieces composed by 483 composers from different historical periods and styles,

ranging Baroque, Classical, and Romantic to Impressionist and Modern. There are 2,869 unique compositions and 5,625 unique pieces and movements. [Figure 2](#) highlights the distributions of pieces and performances per piece for popular composers. [Figure 3](#) shows the distribution of the number of musical pieces by the number of performances. The median and mean numbers of performances per piece are equal to 8 and 44, respectively. In total, 1,104 musical pieces have 50+ performances samples.

Note that **PianoCoRe-C** is not deduplicated or filtered for quality. This raw, comprehensive collection serves as the foundation for the refined subsets, **PianoCoRe-B** and **PianoCoRe-A**, detailed next.

3.6.1 Content and metadata

All score and performance files are organized under the `composer/composition/movement/` directory hierarchy, making the dataset easy to navigate and parse. The following unified naming convention is used:

- **Composer:** composer directories follow IMSLP format `[last_name],[_first_name];`
- **Piece:** piece/opus numbers are represented using Arabic numbers, scales follow the format `[Note]_[?sharp|flat]_[major|minor];`
- **Filename:** The source of every file is preserved in the metadata and the filename, formatted as `[source]_[original_filename].mid.`

The dataset provides content and metadata to support various performance analysis and modeling tasks:

- **Score:** MusicXML and MIDI files, source (ASAP, ATEPP, PDMX, or MuseScore), note count;
- **Performance:** MIDI file, source (ASAP, ATEPP, GiantMIDI, PERiScoPe, or Aria), flag for a transcribed performance and transcription model name,




Figure 2 Statistical overview of the **PianoCoRe-C** dataset for the 50 most represented composers. **Top:** The total number of unique pieces per composer (blue) and the number of pieces with a musical score (light blue). **Bottom:** The average number of performances per piece, accumulated by the MIDI source.

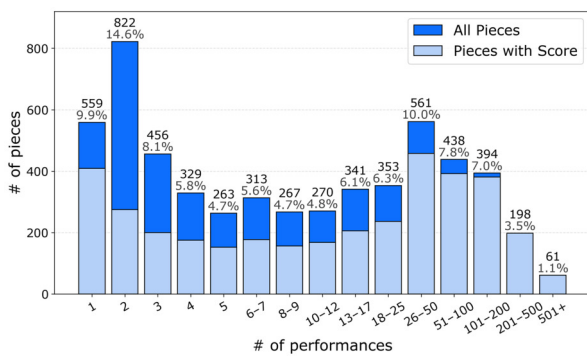


Figure 3 Distribution of the number of musical pieces by the number of performances in **PianoCoRe-C**.

performer’s name (if available), duration and note count;

- **Quality Labels:** lead performance (the higher-priority version of the performance for duplicates), MIDI quality class probabilities and predicted label (‘score’, ‘high quality’, ‘low quality’, ‘corrupted’) (Section 4);
- **Alignment:** if available, path to the `_align.npz` file with raw alignment (after Parangonar), path to the `_refined_align.npz` file with the complete note-to-note alignment between the score and cleaned performance, and alignment recall/precision before and after alignment refinement (Section 5); and
- **Refined Performance:** if alignment is available, refined MIDI file (real and synthetic notes annotated using MIDI markers) that has a complete note alignment with the score MIDI file (Section 5).

3.6.2 Applications

PianoCoRe-C includes matched score and performance MIDI files from the existing piano score and performance datasets: ASAP, ATEPP, PDMX, GiantMIDI-Piano, Aria-MIDI, and PERiScoPe. The combined dataset can be used for tasks that benefit from maximum data scale, such as self-supervised pre-training of music models, large-scale music analysis, or developing data cleaning and filtering techniques.

4 PERFORMANCE MIDI QUALITY ASSESSMENT

The **PianoCoRe-C** dataset contains MIDI files of varying quality, including duplicates. This limits its application to expressive performance modeling. This section details the two-stage refinement process used to produce **PianoCoRe-B** (‘B’ for ‘Base’), a deduplicated, and quality-labeled subset of the data.

4.1 CONTENT-BASED PERFORMANCE DEDUPLICATION

The dataset combines transcribed piano performance MIDI files from multiple sources. The same performance could appear multiple times, either transcribed by different models or uploaded originally under different titles. Duplicates do not add new information and distort the performance data distribution.

For each piece in **PianoCoRe-C** with multiple performances, the performances are compared pairwise using a content-based heuristic developed to detect and

cluster identical or nearly identical performances based on close note onsets. Steps are as follows:

1. **Note Representation:** For each MIDI performance, extract all notes, sort them by time, shift timings so the first note starts at zero, and group notes by pitch number.
2. **Pairwise Similarity:** Take two performances x and z . For each note x_i in x with pitch $p_i = p$, find the closest by onset time, matching note z_j in z with the same pitch $p_j = p$. Then, count the number of note pairs whose absolute time difference is below a threshold $\Delta o_{ij} = |o_i(x) - o_j(z)| \leq 0.05$ (50 ms, an error bound for a near-perfect onset prediction accuracy in AMT (Kong et al., 2021)). The similarity score is the ratio of close note pairs to the total number of notes in x . This score was computed in both directions, from x to z and from z to x , and the maximum was taken.
3. **Clustering:** Performances with at least 50% similar (close in time) notes are clustered. One ‘lead’ file is kept, prioritizing the source datasets with fewer performance samples (GiantMIDI → ATEPP → PERiScape → Aria-MIDI) and, when available, alignment recall.

Applying this method flagged 34,452 near-duplicates, which were removed from the **PianoCoRe-C** dataset, leaving only lead and unique performances. The duplicates are marked in the metadata.

4.2 MIDI QUALITY ASSESSMENT

Besides duplicates, MIDI files transcribed from audio can vary in quality. Since transcription models are trained on limited ground-truth data, they often fail in unseen acoustic conditions (Edwards et al., 2024; Hu et al., 2024). While prior work has proposed perceptually validated metrics (Simonetta et al., 2022; Ycart et al., 2020) and

analytical tools (Hu et al., 2024) for evaluating transcriptions, these methods are reference-based and require ground-truth data for comparison.

Heuristics such as NOMML (Lee et al., 2025) have been used to detect inexpressive MIDI data, but they can struggle with transcriptions. In the experiments, NOMML flagged only 29 performances in PianoCoRe as inexpressive. Transcription artifacts, such as onset jitter, create enough variation to mask a constant tempo, causing score-like performances to appear expressive.

Not all source MIDI performances in PianoCoRe have corresponding audio or musical scores. To classify each performance, a classifier that assesses MIDI quality directly, independent of score and audio alignment, is trained. The main goal is to detect **corrupted** transcriptions and **score-like** performances transcribed from audio synthesizing inexpressive scores.

4.2.1 Note alignment and MIDI quality

The initial hypothesis is that a proxy for MIDI quality is its alignment with the score. The analysis began by examining the differences between recorded performances in ASAP and transcribed performances in ATEPP. In ATEPP, 28.3% of sequences are labeled as ‘high quality,’ ‘low quality,’ ‘background noise,’ or ‘corrupted.’ Figure 4 visualizes the performances using the note ratio $R_n = N_p/N_s$ and adjusted alignment ratio $R'_a = \max(R_n, P_n)$ (Section 3.1). This formulation rewards performances that fully align with the score, even if some segments are not performed.

As we see in Figure 4, ‘recorded’ and ‘high quality’ performances cluster in the upper part ($R'_a > 0.85$), indicating strong alignment with the scores. In contrast, ‘corrupted’ files are inconsistently scattered, including both well- and poorly-aligned performances, while ‘low quality’ and ‘background noise’ sequences overlap with high-quality and corrupted transcriptions.




Figure 4 MIDI performances from ASAP (orange) and ATEPP (blue) grouped by original labels and mapped as a function of performance-to-score note ratio R_n and adjusted alignment ratio R'_a .

Manual inspection of the MIDI files revealed inconsistencies in the original ATEPP labels. Some ‘low quality’ and ‘unlabeled’ files with poor alignment (e.g., 02709.mid, 03001.mid, 10193.mid) contain clearly broken transcriptions. In contrast, a few files labeled as ‘corrupted’ (e.g., 01591.mid, 05389.mid) align well and are musically usable. Thus, the existing audio-based labels do not reliably reflect MIDI quality.

4.2.2 MIDI quality training dataset

Based on the analysis of alignments and the adjusted ratio R'_o , a soft data-labeling heuristic is proposed. Combined with score and recorded MIDI files the four quality classes are defined as follows:

1. **Score (S):** deadpan score MIDI performances;
2. **High Quality (HQ):** any recorded MIDI, transcribed MIDI with $R'_o > 0.9$;
3. **Low Quality (LQ):** transcribed, $0.7 < R'_o < 0.85$; and
4. **Corrupted (C):** transcribed, $R'_o < 0.65$.

The quality ranges are chosen to be disjoint at the boundaries to create clearer distributions for training.

The heuristic was applied to label the deduplicated performances aligned with musical scores. Table 2 shows the distribution of the soft quality labels.

These data were used to sample subsets for training, testing, and calibration. To ensure composition leakage, a piece-based split was applied, maximizing the number of the real corrupted samples in the test set. Second, to create a diverse dataset, there are no more than three samples for each musical piece from each data source (ASAP, ATEPP, PERiSCOPE, and Aria-MIDI), as well as a soft quality label (HQ, LQ, and C).

As seen in Table 2, LQ and C soft labels are underrepresented. For training, 2,500, 1,000, and 86 real HQ, LQ, and C samples, respectively, are balanced with synthetic performances built from the sampled HQ MIDI files. The artificial corruptions for LQ/C classes included random note removal (15%–25%/35%–50%), onset/offset jitter (up to 20 ms/150 ms), velocity jitter (up to 5/20 bins), and random note insertions (up to 5%/30%). Similarly, 953 real scores were augmented with 1,447 synthetic versions (randomized constant velocities, 10-ms onset jitter) to simulate transcription artifacts for score-based audio.

For validation and testing, 200 real Score HQ, and LQ samples are selected alongside 54 Corrupted performances. The classifier calibration set includes all of the real samples from the evaluation split, with no more than three samples per piece, source, and class. The class distributions per each set are shown in Table 3.

4.2.3 MIDI quality classifier

The data representation consists of a stacked sequence encoding with five note features: Pitch, TimeShift (s),

HQ	LQ	C	No Label
170,312	4,545	140	40,597

Table 2 Distribution of MIDI quality labels computed using the alignment-based heuristics for the deduplicated, aligned performances in PianoCoRe-B.

	S	HQ	LQ	C
training	2,500	2,500	2,500	2,500
real	953	2,500	1,000	86
synth	1,547	0	1,500	2,414
test	200	200	200	54
calibration	662	6,525	893	54

Table 3 MIDI quality classification dataset splits.

Velocity (MIDI bins), Duration (s), and absolute TimePosition (s). This encoding does not contain any score features (beat positions and durations) to make the model score-agnostic and universal.

The backbone is a 12-layer transformer encoder (Vaswani et al., 2017) with 80 million parameters, pre-trained using a multi-mask language modeling objective (Borovik et al., 2025). The model dimension is set to 768, and self-attention is extended with Rotary positional embeddings (Su et al., 2024). Real-valued note features are passed to sinusoidal embeddings (Guo et al., 2023) for lossless encoding. For classification, penultimate-layer embeddings are prepended with a [CLS] token and processed by a one-layer transformer (dimension 128) and a classification head.

The pre-training was conducted on the deduped subset of Aria-MIDI (Bradshaw and Colton, 2025) with 371,053 diverse piano MIDI files, provided with the official dataset release. The maximum context length is set to 512 notes. The pre-training included 600,000 steps with batch size 128, while the fine-tuning took 20,000 steps with batch size 512. Training data augmentation included pitch shift (± 6 semitones), velocity shift (± 6 MIDI bins), and tempo stretching ($\pm 5\%$).

The trained MLM backbone was verified on emotion and pianist classification tasks. On the EMOPIA dataset (Hung et al., 2021), the classifier achieved a test accuracy of 72.7% and an F1 score of 72.1%. On the Pianist8 dataset (Chou et al., 2024), the accuracy and F1 score were 86.4% and 85.5%. The metrics are close to similarly sized models (Liang et al., 2024) and slightly below those of larger models (Bradshaw et al., 2025).

4.2.4 Results

Table 4 shows the evaluation results of the classifier configurations tested on the balanced test set.

Model	S	HQ	LQ	C	Avg.
base	1.000	0.839	0.777	0.946	0.891
no synth	1.000	0.759	0.778	0.946	0.871
mean	1.000	0.828	0.752	0.881	0.865
mean, no TL	0.993	0.802	0.713	0.851	0.840
no MLM	0.995	0.773	0.667	0.842	0.819
mask Pitch	1.000	0.803	0.723	0.913	0.860
mask Timing	0.990	0.788	0.747	0.851	0.844
mask Velocity	1.000	0.834	0.776	0.893	0.876

Table 4 Evaluation of MIDI quality classifiers using F1 scores. Best scores in **bold**. no synth—no synthetic training data, mean—mean pooling (no [CLS]), no TL—no transformer layer before the classifier head, no MLM—token embeddings and classifier only. The last block shows feature-masking ablations.

The best configuration achieved a macro F1 score of 89.1% on the held-out test set. It learned to perfectly distinguish score-like MIDI files and showed less errors between HQ, LQ, and C classes. The synthetic training samples and token-based aggregation helped to learn more robust decision boundaries. Masking of note features revealed the shared contribution of pitch, dynamic, and timing to MIDI quality classification. Since note-level alignment is imperfect and quality is continuous rather than discrete, errors on the test set are expected.

4.3 CLASSIFYING THE PianoCoRe-C DATASET

The best-performing classifier was taken and calibrated on the held-out calibration set (Table 3). To maximize recall, the sequences are labeled as Corrupted or Score, if the classifier was activated in at least one segment ($p_s > 0.3$ or $p_c > 0.3$). For the LQ class, a conservative threshold of $p_{LQ} > 0.75$, which does not categorize half of the data as low quality, was chosen. Note that HQ and LQ labels are advisory, as ‘low quality’ MIDI files may be suitable for certain applications. However, the files labeled as Corrupted or Score are, in most cases, indeed either broken or were transcribed from rendered musical scores with constant tempo and/or dynamics. It is better to filter them during piano-expression analysis.

The final distribution of MIDI quality labels in the PianoCoRe-C dataset is shown in Table 5.

4.4 PianoCoRe-B DATASET

By applying the deduplication and quality assessment models to PianoCoRe-C dataset, we obtain PianoCoRe-B. The filtered subset consists of 214,092 deduplicated performance MIDI not classified as Corrupted or Score. There are 5,591 musical pieces composed by 478 composers (Table 1).

Source	S	HQ	LQ	C
ASAP	0	1,066	0	0
ATEPP	0	10,231	900	433
GiantMIDI	11	2,071	52	5
PERIScoPe	82	34,596	91	4
Aria-MIDI	1,151	180,977	18,359	17
PianoCoRe	1,244	228,941	19,402	459

Table 5 PianoCoRe dataset and its source subsets labeled by the MIDI quality classifier.

4.4.1 Applications

PianoCoRe-B is designed for tasks that depend on large amounts of clean and reliable piano performance data. Specifically, this dataset is useful for large-scale, self-supervised pre-training; musical analysis of performance styles; and piano performance generation.

5 REFINED NOTE ALIGNMENT

Piano-expression modeling tasks require precise note-level alignment between scores and performances. The **PianoCoRe-A/A*** subsets (‘A’ for ‘Aligned’) consist of all performance MIDI files that are temporally aligned to scores. Two forms of alignment are considered:

- 1. Raw Alignments:** processed output of Parangonar, containing matches, insertions, and deletions between score and performance notes and
- 2. Refined Alignments:** raw alignments, refined using the **RAScoP** pipeline, which cleans and completes initial matches.

5.1 RAW NOTE ALIGNMENT CHALLENGES

A direct output from note aligners like Parangonar (Peter, 2023) or Nakamura’s alignment tool (Nakamura et al., 2017), while powerful, is sometimes insufficient for direct use in generative models. Raw alignments can suffer from issues, illustrated in Figure 5:

- **Temporal Discontinuities:** Incorrect alignment links that cross in time or match musically distant notes, leading to unrealistic tempo fluctuations and high inter-onset timing deviations;
- **Alignment Holes:** Continuous regions of unaligned notes in the score or performance, often caused by skipped repeats or transcription errors.

Some performance rendering models were trained only on a subset of aligned score and performance notes with incomplete score contexts (Rhyu et al., 2022; Tang et al., 2025; Zhang et al., 2024). Other models removed

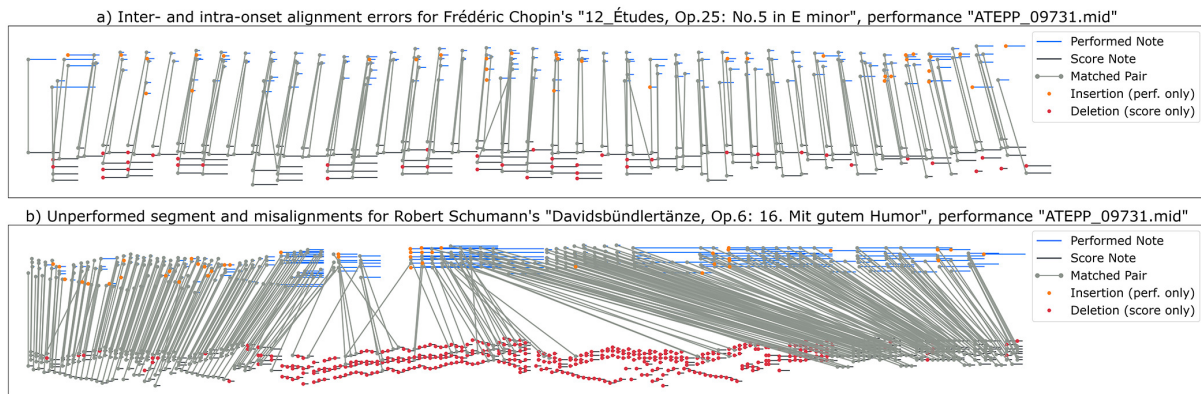


Figure 5 Real-world alignment challenges motivating the RAScoP pipeline. **Top:** local timing errors (crossed links) and missing/extra notes. **Bottom:** large structural deviation from a missing score segment, causing incorrect links. Other performed notes remain usable. Alignments were computed with Parangonar.

timing outliers (Jeong et al., 2019a; Xia, 2016) and interpolated missing notes (Borovik and Viro, 2023; Borovik et al., 2025). However, these processes are not available as easy-to-use tools.

A configurable algorithm was designed to create a parallel score and performance dataset by cleaning evident outliers and interpolating notes for which no performance counterpart exists. Specifically, this algorithm addresses two main problems:

- **Timing Errors:** remove large inter- and intra-onset deviations and implied unrealistic tempi and
- **Missing Notes:** fill in the unperformed notes to have complete performed score contexts.

The following section describes this algorithm.

5.2 ALIGNMENT CLEANING AND REFINEMENT

RAScoP ('Refined Alignment for Scores and Performances') is an integrated pipeline designed to take a raw score-performance alignment and transform it into a clean, complete, and temporally coherent parallel score-performance data pair. The algorithm analyzes and refines the alignment through four sequential steps, illustrated in Figure 6:

1. **(H):** alignment hole processing,
2. **(O):** onset cleaning and temporal refinement,
3. **(I):** note interpolation, and
4. **(S):** performance-to-score synchronization.

5.2.1 Alignment hole processing

The first step detects and removes large, structurally incorrect alignment sections. An 'alignment hole' is defined as a continuous region of notes where the alignment is sparse or nonsensical (only a few notes are aligned). In scores, the holes correspond to unperformed score measures (e.g., repeats), whose individual notes may be incorrectly matched with random performance

notes. In the performances, the holes are the extra performed segments whose notes may be inadvertently aligned with random score notes.

To detect holes, a sliding window approach is used. Let H_o be a ratio of unaligned notes within a surrounding window of H_w notes for a given note. If H_o ratio exceeds a threshold H_r , the note is flagged. Contiguous regions of flagged notes are designated as holes, and all alignment pairs within them are removed.

The default values are $H_w = 31$ notes and $H_r = 0.75$. The window size is close to double the median (15) and mean (16.9) number of notes in a measure in all scores in the dataset. With this window, we consider on average one measure to the left and one to the right. Setting the threshold at 75% ensures that only regions that are almost entirely unaligned are removed.

5.2.2 Onset cleaning and temporal refinement

This stage refines the temporal alignment of concurrently played notes (chords) and corrects large-scale time shifts. First, all aligned notes are used to build the initial onset pair list: tuples of score onset beat o_i and the average performed onset time $t(o_i)$ for all notes in the chord. Then, note and onset times are checked for misalignments and outliers based on:

- high intra-onset deviations and
- inter-onset intervals that deviate from the local performance tempo.

For intra-onset deviations, the onset deviations $\Delta t_i(n_j) = t(n_j) - t(o_i)$ from $t(o_i)$ are computed for all notes n_j in a chord: $\{n_j | o(n_j) = o_i\}$. By default, notes whose onsets deviate from $t(o_i)$ by more than two standard deviations are removed from the alignment as outliers. For chords with two distant notes, both notes will be removed if the condition is met.

For inter-onset intervals, the method estimates the maximum and minimum plausible time shifts $\Delta t_{\max}(o_i)$

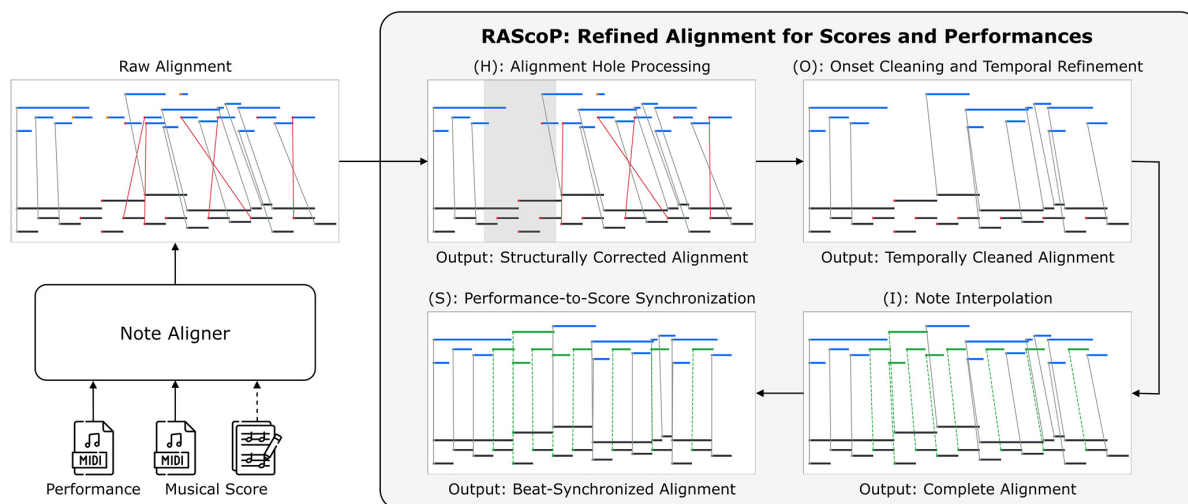


Figure 6 Note-level alignment and the RASCoP pipeline for alignment refinement. The processing steps are demonstrated using an artificial example containing all types of errors. Score notes are drawn in black and performance notes are drawn in blue and green.

and $\Delta t_{\min}(o_i)$ between the current and previous score onsets o_i and o_{i-1} . If the time interval between the current and previous onset implies a tempo outside a plausible range (by default, 15–480 BPM), it is identified as an alignment jump. This onset can be filtered out of the alignment. However, by default, the timing of the notes of the affected onsets is adjusted.

First, a local tempo $\tau_{\text{local}}(o_i)$ is estimated based on a w -second window (by default, $w = 8$) of preceding performed note onsets $O_{\text{local}}(o_i) = \{o_j | t(o_j) - t(o_i) < w\}$. Then, the expected onset time $\hat{t}(o_i)$ is computed using the inter-onset beat shift $\text{IOI}_i^s = o_i - o_{i-1}$ and $\tau_{\text{local}}(o_i)$. Using the expected onset time, the required time shift $\Delta t_{\text{adj}}(o_i) = \hat{t}(o_i) - t(o_i)$ is determined, and the subsequent performance notes are shifted accordingly.

This step explicitly alters the global timing in the original performance MIDI. However, after the shift, the onsets fall into the range of the plausible local tempos, and tempo outliers are not learned by the trained models. Any unperformed notes can be also naturally filled in with the same local performance tempo.

In addition, close onset pairs with $\Delta t(o_i) = t(o_i) - t(o_{i-1}) < 0.01$ (10 ms) are filtered out to avoid two same-pitch note-on events (which is impossible for a human performer). After the alignment hole processing and onset cleaning, the algorithm cleans up the performance MIDI by removing notes without a link in the alignment. In the end, only matched and cleanly performed notes remain.

5.2.3 Note interpolation

This step interpolates the unperformed notes to create parallel note-aligned score–performance pairs.

The note onset time $t(n_i)$ of a note n_i is linearly interpolated from two neighboring performed notes n_j and n_k . To avoid the contribution of very close notes, the configurable minimum beat and time intervals n_j and n_k

between the two anchor notes are used ($t(n_k) - t(n_j) \geq \Delta t_{\text{int}}$ and $o(n_k) - o(n_j) \geq \Delta o_{\text{int}}$).

Note articulation (duration) and dynamics (MIDI velocity) are averaged and weighted by the performed notes in the neighboring beats. The weights are inversely proportional to the absolute beat distances $\text{IOI}_{i,j}^s = |o(n_j) - o(n_i)|$ from the score position of the note n_i being interpolated. Closer notes contribute a higher weight to the interpolated features.

The algorithm prevents the creation of notes with identical pitch and onset, and shortens overlapping notes so that at each new key press the previous note is closed. The result is a performance MIDI file fully aligned with the score at the note level. Interpolated notes are marked with a special MIDI text marker, allowing them to be filtered out or marked during model training.

5.2.4 Performance–score synchronization

This step synchronizes the beat structure of the refined performance MIDI with the score. This data format is commonly used in MIDI encodings with beat/bar tempo (Huang and Yang, 2020; Hsiao et al., 2021; Zeng et al., 2021). The alignment pairs are used to compute a beat-to-time mapping and insert inter-beat tempo changes into the performance MIDI. For example, for a 4/4 time signature and 480 ticks per quarter, notes at the beats are separated by 480 ticks in both the score and performance MIDI, with exact times derived from tempo changes.

Finally, the entire performance is shifted so that its first played note occurs at the same time as the first score note, ensuring a consistent starting point for all performances of the same composition.

5.2.5 Final output

The algorithm returns the refined alignment, refined performance MIDI, and note-level alignment recall ratios. The recall values from different stages (initial, hole

processing, and onset cleaning) serve as quantitative indicators of alignment quality and can be used to interrupt the refinement process. The alignment is released as a compressed .npz file containing an array of performance note indices aligned to the sorted score MIDI notes, along with a boolean mask for interpolated notes. All MIDI processing steps are performed using the `symusic` Python library (Liao et al., 2024).

The presented refinement does not rematch links produced by the note aligner. It only filters existing links and interpolates missing notes. Each step of the pipeline can be enabled independently. Default parameters were chosen empirically rather than optimized, as automated evaluation would require precise human annotations. Custom clean datasets can be generated using the released raw alignments and MIDI files.

In PianoCoRe, all refined performance MIDI files underwent the first three stages of the RAScoP alignment and the final initial performance onset shift. Beat synchronization was not applied in order to preserve the original timing without re-quantizing the note onsets and offsets. Synchronization can be computed using the refined score MIDI, performance MIDI, and note-level alignment.

5.3 REFINEMENT QUALITY EVALUATION

To quantitatively demonstrate the effectiveness of RAScoP, the trade-off between alignment temporal integrity (the distribution of intra- and inter-onset deviations) and alignment recall (R_o) is evaluated.

The benefit of alignment refinement is shown in Figure 7. Applying the full pipeline (H + O) significantly reduces the standard deviation of inter-onset deviations within chords, indicating cleaner note timing patterns. Furthermore, the distribution of beat tempos becomes more stable and centered around a musically plausible range, as the algorithm corrects for the extreme tempo values implied by raw, noisy alignments.

Table 6 quantifies the ‘cost’ of the cleaning process in terms of alignment recall. The performances are grouped from higher to lower recall. Overall, the average recall \bar{R}_o decreases by a modest 1.5% (from 0.935 to 0.920), with the Onset Cleaning stage (O) contributing most to this reduction. The cleaning process primarily affects the highest-quality alignments ($\bar{R}_o > 0.95$), reducing their share from 54.3% to 42.9%. These sequences are not discarded but rather migrate to the still-high-quality lower bands. After refinement, the majority of sequences (86.6%) still maintain a high alignment recall of over 85%. The loss of a few alignment links is an acceptable price for the improvement in the temporal quality of the performance data.

5.4 PianoCoRe-A DATASET

Applying this pipeline to the files from PianoCoRe-B yields the final aligned datasets. PianoCoRe-A contains 157,207 cleaned and note-aligned sequences from

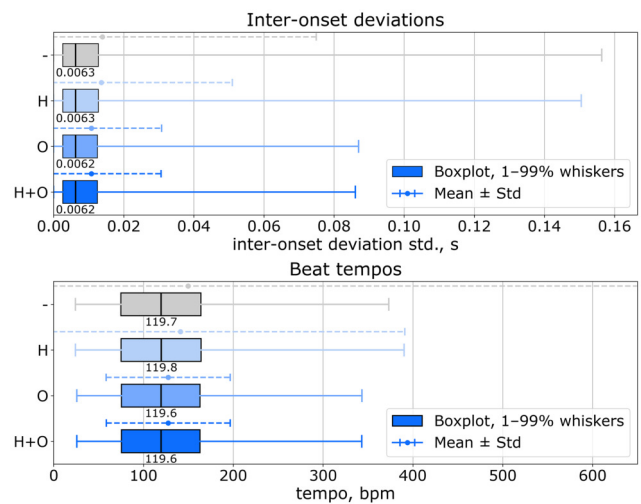


Figure 7 Distribution of inter-onset deviations and beat tempos for alignments before processing (-), after hole processing (H), after onset cleaning (O), and after both hole and onset cleaning (H + O).

Band	Raw		After H		After H+O	
	\bar{R}_A	%	\bar{R}_H	%	\bar{R}_{H+O}	%
0.95–1.00	0.975	54.3	0.975	53.9	0.973	42.9
0.90–0.95	0.929	26.6	0.929	26.7	0.928	30.4
0.85–0.90	0.879	10.1	0.878	10.0	0.878	13.3
0.80–0.85	0.828	4.7	0.828	4.6	0.828	6.5
0.75–0.80	0.779	2.1	0.778	2.2	0.777	3.2
0.70–0.75	0.725	1.1	0.727	1.0	0.728	1.6
0.60–0.70	0.660	0.7	0.663	1.1	0.661	1.5
0.00–0.60	0.471	0.4	0.464	0.5	0.462	0.6
all	0.935	100.0	0.934	100.0	0.920	100.0

Table 6 Mean alignment recall \bar{R} after different alignment refinement stages and the ratio of sequences (%) inside different recall bands.

PianoCoRe-B for 1,591 pieces written by 151 composers, totaling 12,509 hours of music (Table 1).

The performances can be filtered out for any applications based on the alignment ratio. For tasks that demand the highest possible data fidelity, PianoCoRe-A* is introduced. This is a high-confidence subset of PianoCoRe-A containing High Quality MIDI files with at least 85% of aligned notes. PianoCoRe-A* consists of 130,275 performances for 1,517 pieces.

5.4.1 Applications

PianoCoRe-A/A* represent a large-scale resource of score–performance–aligned piano MIDI data. They pave the way for training more nuanced models for rendering expressive piano performances without having to perform rigorous data matching and alignment.

6 MUSIC PERFORMANCE RENDERING

The PianoCoRe dataset is validated on a downstream task of expressive piano performance rendering. The hypothesis is that the scale, diversity, and targeted refinement of the **PianoCoRe-A** dataset enable the training of more accurate performance models compared to baselines trained on smaller or uncleaned data subsets.

6.1 EXPERIMENTAL SETUP

The experiments used PianoFlow (Borovik et al., 2025), a model for symbolic music performance rendering based on conditional flow matching (Lipman et al., 2022). It employs an encoder transformer to inpaint masked performance features x_m (TimeShift, Velocity, TimeDuration, TimeDurationSustain) given score features y (Pitch, Position, PositionShift, and Duration) and performance context x_{ctx} . As Aria-MIDI does not distinguish between pressed and sustained notes, only seven features without TimeDuration were used. The base configuration (8 layers, 24 million parameters) was adopted, and a learned embedding was added to interpolated notes, as in the original model.

The model was trained on subsets of aligned and cleaned performances from PianoCoRe-A: ASAP, ASAP+ATEPP, ASAP+ATEPP+PERiScoPe, and the full dataset. Performances with fewer than 85% aligned notes ($R_{RASCoP} < 0.85$) were removed to retain more real played notes. For ablation, models were trained on all PianoCoRe-A performances ($R_{RASCoP} \geq 0.7$) and a version of the dataset without the hole and onset cleaning from RASCoP pipeline (raw alignments plus note interpolation). Data were split by composition into 90%/10% for training/evaluation, all movements and performances of a piece appeared in only one split.

6.2 RESULTS

6.2.1 Training convergence

Figure 8 illustrates the feature-based validation losses tracked during training. Each model was evaluated on a validation set drawn from the same source data (e.g., the ‘ASAP+ATEPP’ model on unseen ‘ASAP+ATEPP’

performances). The results reveal a pattern: the model trained only on ‘ASAP’ quickly overfits, demonstrating that a small dataset, even of high quality, is insufficient. As the scale of the data increases (‘+ATEPP’, ‘+PERiScoPe’), overfitting is delayed.

The comparison between the ‘PianoCoRe-A’ model (blue) and its unrefined counterpart ‘w/o RASCoP’ (gray) provides direct evidence of the value of the refinement pipeline. The refined dataset yields a more stable and consistently lower validation loss, particularly for the note time shifts. This confirms that targeted removal of temporal noise is crucial for learning an accurate timing model.

6.2.2 Unconditional generation

This section presents the evaluation results for the unconditional performance rendering. The inference set included test set scores with at least three performances from two different MIDI sources (e.g., ASAP and Aria-MIDI). The models rendered each score in its entirety seven times. Pearson correlation (Borovik and Viro, 2023; Jeong et al., 2019b; Zhang et al., 2024) between the note features of the dataset and rendered performances was computed. The evaluated features are: onset velocity (Vel), relative inter-onset intervals (IOI), relative intra-onset deviations (OD), and note articulation (Art).

Table 7 presents the mean Pearson correlation between the model outputs and the ground-truth performances from a multi-source test set. Models trained on more diverse datasets (‘+ ATEPP’, ‘+ PERiScoPe’, and ‘PianoCoRe-A’) consistently outperform the baseline trained only on ‘ASAP’. Interestingly, the model trained on ASAP and ATEPP shows higher correlation with an average set of performances from PianoCoRe-A. This may be because ATEPP specifically focuses on the performances of renowned pianists. Other datasets contain a wider variety of performance styles.

More training data with more interpolated notes ($R_{RASCoP} \geq 0.7$) slightly hurts the unconditional rendering capabilities. The model trained on raw data without the cleanup shows lower correlation with higher quality performances for note timing (IOI and OD).




Figure 8 Validation loss curves for PianoFlow trained on different subsets of the data. Larger and refined training datasets reduce overfitting in the long run.

6.2.3 Performance continuation

The final analysis evaluated the models in a performance continuation task across four distinct test domains: ASAP, ATEPP, PERiScoPe, and Aria. As in the previous experiments, compositions and performances were not seen during the training. The models performed 256 notes in parallel using the performance context of the preceding 256 notes. Table 8 shows the mean absolute error computed against the ground truth performance features.

The results complement the previous findings. With more training data, the model performs better on MIDI files of different sources. PianoCoRe-A achieves the best average performance on ASAP and Aria-MIDI performances and second-best results on the other subsets. Only the model trained on data without overrepresented Aria-MIDI achieves similar or lower errors on ATEPP and PERiScoPe. Given the validation loss plots in Figure 8, the full dataset model has room for an improvement in the long run. Overall, the results show the potential

	Vel	IOI	OD	Art
Dataset	0.57±0.19	0.90±0.06	0.22±0.17	0.44±0.19
ASAP	0.37±0.17	0.83±0.11	0.07±0.15	0.28±0.13
+ ATEPP	0.42±0.16	0.85±0.11	0.12±0.14	0.35±0.15
+ PERiScoPe	0.41±0.17	0.86±0.11	0.11±0.17	0.36±0.17
PianoCoRe-A	0.40±0.17	0.86±0.11	0.10±0.17	0.35±0.17
$R_{\text{RASCoP}} \geq 0.7$	0.39±0.16	0.85±0.11	0.09±0.16	0.35±0.18
w/o RASCoP	0.41±0.16	0.85±0.11	0.09±0.16	0.36±0.18

Table 7 Correlation between the features of the rendered and PianoCoRe-A performances. First row—*intra-set* correlations, other rows—models trained on different data subsets. *Vel*—velocity, *IOI*—inter-onset-interval, *OD*—relative onset deviation, *Art*—sustained articulation. The best scores are in **bold**.

Dataset	Size	ASAP			ATEPP			PERiScoPe			Aria-MIDI		
		Vel	TS	TD	Vel	TS	TD	Vel	TS	TD	Vel	TS	TD
ASAP	1 k	9.885	0.023	0.187	9.928	0.022	0.206	9.893	0.023	0.230	9.957	0.027	0.275
+ ATEPP	6 k	9.157	0.017	0.168	8.230	0.015	0.191	8.782	0.016	0.216	8.721	0.019	0.252
+ PERiScoPe	25 k	8.851	0.016	0.154	7.888	0.013	0.189	8.117	0.015	0.192	8.133	0.017	0.230
PianoCoRe-A	124 k	8.613	0.016	0.155	7.967	0.014	0.194	8.094	0.015	0.194	7.872	0.017	0.205
$R_{\text{RASCoP}} \geq 0.7$	141 k	8.631	0.016	0.158	7.944	0.014	0.196	8.071	0.015	0.194	7.921	0.017	0.206
w/o RASCoP	124 k	8.734	0.017	0.159	8.059	0.015	0.193	8.199	0.016	0.196	8.055	0.018	0.211

Table 8 Conditional performance rendering (performance continuation) results across training subsets and unseen source sequences. *Size* denotes the training set size. *Vel*—Velocity (MIDI bins), *TS*—TimeShift (s), *TD*—TimeDurationSustain (s). Lower is better; best values are in **bold**.

of PianoCoRe for training performance models robust to varying piano data distributions.

6.3 FUTURE WORK

A subjective listening test of popular models trained on the subsets of PianoCoRe dataset would be a valuable next step to confirm that objective improvements translate to human perception. Since performances from Aria-MIDI dominate PianoCoRe, a more balanced sampling of performances per source might provide a better generalization to all source data domains. Fine-tuning on high-fidelity subsets, such as ASAP, could potentially improve performance even further.

7 LIMITATIONS

Despite rigorous curation, PianoCoRe has limitations. There are no duplicate musical pieces with different names. However, an error margin of 1% is reserved for potential movement-level naming errors that were inherited from the source datasets. Furthermore, the dataset distribution remains skewed toward Western classical repertoire and popular composers, reflecting the biases of the underlying open-source corpora.

The dataset relies on open-source MusicXML scores and automated alignment. MusicXML scores are not error-free and may also include a segment of a complete written musical composition. Since it is difficult to validate large-scale datasets precisely, any errors in the source notations may propagate to the downstream applications. Also, due to the iterative combination of source datasets, fewer than 1% of performances may contain neighboring movements or differ from the scores by more than twice the length. It is recommended to use composition-wise splits in the applications using the dataset.

The classifier-based MIDI quality labels were calibrated for recall in the corrupted and score-like classes to filter out incorrect and inexpressive data. The labels do not guarantee perfect alignment with human expectations. During note interpolation, RAScoP may introduce deadpan performance note segments that must be addressed by downstream applications. Additionally, interpolation does not handle sustain pedal effects. A better solution would be to predict missing notes and pedals using a trained model.

8 CONCLUSION

This article presented **PianoCoRe**, a unified, large-scale piano MIDI dataset created by combining, refining, annotating, and aligning existing open-source corpora. Released in tiered subsets, PianoCoRe supports a wide spectrum of tasks: from performance analysis and large-scale pre-training to expressive piano performance rendering and score-to-performance translation. The dataset enables reproducible research by allowing researchers to create non-overlapping data splits across previously isolated datasets.

To ensure data integrity, two challenges were addressed: the quality of performance MIDI and note-level alignments. A classifier was trained to identify deadpan and corrupted MIDI transcriptions, and an alignment refinement pipeline was designed to remove temporal outliers in aligned score-performance data. The experiments showed that the model trained on these refined subsets benefits from the increased repertoire diversity and cleaner note features.

Future directions include extending the methodology to multi-instrument repertoires, developing more robust quality assessment models and incorporating more granular score and performance annotations. By making PianoCoRe openly available, the goal is to establish a foundation for advancing symbolic music performance modeling and analysis research.

ACKNOWLEDGMENTS

The author would like to thank Vladimir Viro and Dmitrii Gavrilov for their feedback and suggestions regarding early versions of the alignment refinement algorithm and the dataset. The author is grateful to the TISMIR editorial team and the anonymous reviewers for their constructive and invaluable feedback, which improved the quality of the dataset and manuscript.

The work was made possible by the use of the Zhores cluster and its computational resources (Zacharov et al., 2019). Furthermore, the author expresses gratitude to the creators of the MAESTRO, ASAP, (n)ASAP, ATEPP, GiantMIDI-Piano, Aria-MIDI, and PERIScoPe datasets. Their commitment to open science and the sharing of

symbolic music resources provided the essential foundation for this work.

ETHICAL STATEMENT

The curation of large-scale symbolic datasets presents challenges regarding copyright and intellectual property. A best-effort attempt was made to filter PianoCoRe according to European Union public-domain regulations (works whose authors have been deceased for more than 70 years). However, achieving 100% accuracy across thousands of files from diverse sources is inherently difficult. For transparency, the annotated composer metadata is released alongside the dataset.

The dataset, original and processed files, metadata, and alignment annotations are published under a CC-BY-NC-SA 4.0 license. The license respects the licenses used for the source datasets. No formal ethics approval or human participant consent was required for this study, as it involved the processing of publicly available MIDI data and did not involve human subjects.

DATA ACCESSIBILITY

The PianoCoRe dataset and related resources are released to enforce reproducibility:

- **Code:** The source code, documentation, and usage examples are available at the project repository: <https://github.com/ilya16/PianoCoRe>;
- **Dataset:** The data are archived on Zenodo at <https://doi.org/10.5281/zenodo.19186016> and is available on Hugging Face at <https://huggingface.co/datasets/SyMuPe/PianoCoRe>.

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR'S CONTRIBUTION


Ilya Borovik was responsible for the research conceptualization, methodology, software implementation, data curation, and the writing of the manuscript.

NOTES

1. <https://github.com/ilya16/PianoCoRe>
2. <https://doi.org/10.5281/zenodo.19186016>
3. <https://huggingface.co/datasets/SyMuPe/PianoCoRe>
4. <https://github.com/CPJKU/asap-dataset>
5. <https://github.com/tangjibetsy/ATEPP>
6. <https://github.com/bytedance/GiantMIDI-Piano>
7. <https://huggingface.co/datasets/SyMuPe/PERIScoPe>
8. <https://huggingface.co/datasets/loubb/aria-midi>

9. <https://musescore.com/sheetmusic>
10. <https://kunstderfuge.com>
11. <https://www.classicalmidi.co.uk>
12. <https://imslp.org>
13. <https://eur-lex.europa.eu/EN/legal-content/summary/copyright-and-related-rights-term-of-protection.html>

AUTHOR AFFILIATIONS

Ilya Borovik  <https://orcid.org/0009-0002-1186-2039>
Skolkovo Institute of Science and Technology, Moscow, Russia

REFERENCES

- Benetos, E., Dixon, S., Duan, Z., and Ewert, S.** (2018). Automatic music transcription: An overview. *IEEE Signal Processing Magazine*, 36(1), 20–30.
- Borovik, I., Gavrilov, D., and Viro, V.** (2025). SyMuPe: Affective and controllable symbolic music performance. In *Proceedings of the 33rd ACM International Conference on Multimedia*, Dublin, Ireland, pp. 10699–10708.
- Borovik, I., and Viro, V.** (2023). ScorePerformer: Expressive piano performance rendering with fine-grained control. In *Proceedings of the 24th International Society for Music Information Retrieval Conference (ISMIR)*, Milan, Italy, pp. 588–596.
- Bradshaw, L., and Colton, S.** (2025). Aria-MIDI: A dataset of piano MIDI files for symbolic music modeling. In *Proceedings of the 13th International Conference on Representation Learning (ICLR)*, Singapore, Singapore.
- Bradshaw, L., Fan, H., Spangher, A., Biderman, S., and Colton, S.** (2025). Scaling self-supervised representation learning for symbolic piano performance. In *Proceedings of the 26th International Society for Music Information Retrieval Conference (ISMIR)*, Daejeon, Korea, pp. 451–459.
- Cancino-Chacón, C. E., Grachten, M., Goebel, W., and Widmer, G.** (2018). Computational models of expressive music performance: A comprehensive and critical review. *Frontiers in Digital Humanities*, 5, 25.
- Cancino-Chacón, C. E., Peter, S. D., Karystinaios, E., Foscari, F., Grachten, M., and Widmer, G.** (2022). Partitura: A Python package for symbolic music processing. In *Proceedings of the Music Encoding Conference (MEC)*, Halifax, Canada.
- Chou, Y.-H., Chen, I.-C., Ching, J., Chang, C.-J., and Yang, Y.-H.** (2024). MidiBERT-Piano: Large-scale pre-training for symbolic music classification tasks. *Journal of Creative Music Systems*, 8(1).
- Edwards, D., Dixon, S., and Benetos, E.** (2023). PiJAMA: Piano jazz with automatic MIDI annotations. *Transactions of the International Society for Music Information Retrieval*, 6(1), 89–102.
- Edwards, D., Dixon, S., Benetos, E., Maezawa, A., and Kusaka, Y.** (2024). A data-driven analysis of robust automatic piano transcription. *IEEE Signal Processing Letters*, 31, 681–685.
- Emerson, K., and Harrison, P. M. C.** (2025). Multimodal datasets for studying expert performances of musical scores. *Transactions of the International Society for Music Information Retrieval*, 8(1), 400–428.
- Foscari, F., Mcleod, A., Rigaux, P., Jacquemard, F., and Sakai, M.** (2020). ASAP: A dataset of aligned scores and performances for piano transcription. In *Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR)*, Montréal, Canada, pp. 534–541.
- Goebel, W.** (1999). *The Vienna 4x22 Piano Corpus*. <http://dx.doi.org/10.21939/4X22>.
- Good, M.** (2001). MusicXML for notation and analysis. In *The Virtual Score: Representation, Retrieval, Restoration*, 12, 113–124.
- Guo, Z., Kang, J., and Herremans, D.** (2023). A domain-knowledge-inspired music embedding space and a novel attention mechanism for symbolic music modeling. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, Volume 37, Washington, DC, USA, pp. 5070–5077.
- Hashida, M., Nakamura, E., and Katayose, H.** (2018). Crest-MusePEDB 2nd edition: Music performance database with phrase information. In *Proceedings of the 15th Sound and Music Computing Conference (SMC)*, Limassol, Cyprus.
- Hawthorne, C., Stasyuk, A., Roberts, A., Simon, I., Huang, C.-Z. A., Dieleman, S., Elsen, E., Engel, J., and Eck, D.** (2019). Enabling factorized piano music modeling and generation with the MAESTRO dataset. In *Proceedings of the 7th International Conference on Representation Learning (ICLR)*, New Orleans, LA, USA.
- Hsiao, W.-Y., Liu, J.-Y., Yeh, Y.-C., and Yang, Y.-H.** (2021). Compound word transformer: Learning to compose full-song music over dynamic directed hypergraphs. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence* (Volume 35, pp. 178–186). Virtual Event.
- Hu, P., Marták, L. S., Cancino-Chacón, C., and Widmer, G.** (2024). Towards musically informed evaluation of piano transcription models. In *Proceedings of the 25th International Society for Music Information Retrieval Conference (ISMIR)*, San Francisco, CA, USA, pp. 1068–1075.
- Hu, P., and Widmer, G.** (2023). The Batik-Plays-Mozart corpus: Linking performance to score to musicological annotations. In *Proceedings of the 24th International Society for Music Information Retrieval Conference (ISMIR)*, Milan, Italy, pp. 297–303.
- Huang, Y.-S., and Yang, Y.-H.** (2020). Pop music transformer: Beat-based modeling and generation of expressive pop piano compositions. In *Proceedings of the 28th ACM International Conference on Multimedia*, Virtual Event and Seattle, WA, USA, pp. 1180–1188.
- Hung, H.-T., Ching, J., Doh, S., Kim, N., Nam, J., and Yang, Y.-H.** (2021). EMOPIA: A multi-modal pop piano dataset for emotion recognition and emotion-based music generation. In *Proceedings of the 22nd International Society for Music Information Retrieval Conference (ISMIR)*, pp. 318–325. Online.
- Jeong, D., Kwon, T., Kim, Y., Lee, K., and Nam, J.** (2019a). VirtuosoNet: A hierarchical RNN-based system for modeling

- expressive piano performance. In *Proceedings of the 20th International Society for Music Information Retrieval Conference (ISMIR)*, Delft, Netherlands, pp. 908–915.
- Jeong, D., Kwon, T., Kim, Y., and Nam, J.** (2019b). Graph neural network for music score data and modeling expressive piano performance. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, Long Beach, CA, USA, pp. 3060–3070. PMLR.
- Kong, Q., Li, B., Chen, J., and Wang, Y.** (2022). GiantMIDI-Piano: A large-scale MIDI dataset for classical piano music. *Transactions of the International Society for Music Information Retrieval*, Bengaluru, India, 5(1), 87–98.
- Kong, Q., Li, B., Song, X., Wan, Y., and Wang, Y.** (2021). High-resolution piano transcription with pedals by regressing onset and offset times. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 3707–3717.
- Kosta, K., Bandtlow, O. F., and Chew, E.** (2018). MazurkaBL: Score-aligned loudness, beat, expressive markings data for 2000 Chopin Mazurka recordings. In *Proceedings of the 4th International Conference on Technologies for Music Notation and Representation (TENOR)*, Montréal, Canada, pp. 85–94.
- Lam, S. K., Pitrou, A., and Seibert, S.** (2015). Numba: A LLVM-based Python JIT compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, Austin, TX, USA, pp. 1–6.
- Lee, K. J. M., Ens, J., Adkins, S., Sarmento, P., Barthet, M., and Pasquier, P.** (2025). The GigaMIDI dataset with features for expressive music performance detection. *Transactions of the International Society for Music Information Retrieval*, 8(1), 1–19.
- Lerch, A., Arthur, C., Pati, A., and Gururani, S.** (2020). An interdisciplinary review of music performance analysis. *Transactions of the International Society for Music Information Retrieval*, 3(1), 221–245.
- Liang, X., Zhao, Z., Zeng, W., He, Y., He, F., Wang, Y., and Gao, C.** (2024). PianoBART: Symbolic piano music generation and understanding with large-scale pre-training. In *Proceedings of the 25th IEEE International Conference on Multimedia and Expo (ICME)*, IEEE, Niagara Falls, ON, Canada, pp. 1–6.
- Liao, Y., Luo, Z., Wang, Y., and Yin, Y.** (2024). Symusic: A swift and unified toolkit for symbolic music processing. In *Extended Abstracts of the 25th International Society for Music Information Retrieval Conference (ISMIR)*, San Francisco, CA, USA.
- Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M.** (2022). Flow matching for generative modeling. In *The Proceedings of the 11th International Conference on Learning Representations (ICLR)*, Virtual Event.
- Long, P., Novack, Z., Berg-Kirkpatrick, T., and McAuley, J.** (2025). PDMX: A large-scale public domain MusicXML dataset for symbolic music processing. In *Proceedings of the 50th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, Hyderabad, India, pp. 1–5.
- Müller, M., Konz, V., Bogler, W., and Arifi-Müller, V.** (2011). Saarland Music Data (SMD). In *Extended Abstracts for the Late-Breaking Demo Session of the 12th International Society for Music Information Retrieval Conference (ISMIR)*, Miami, FL, USA.
- Nakamura, E., Yoshii, K., and Katayose, H.** (2017). Performance error detection and post-processing for fast and accurate symbolic music alignment. In *Proceedings of the 18th International Society for Music Information Retrieval Conference (ISMIR)*, Suzhou, China, pp. 347–353.
- Peter, S. D.** (2023). Online symbolic music alignment with offline reinforcement learning. In *Proceedings of the 24th International Society for Music Information Retrieval Conference (ISMIR)*, Milan, Italy, pp. 634–641.
- Peter, S. D., Cancino-Chacón, C. E., Foscarin, F., McLeod, A. P., Henkel, F., Karystinaios, E., and Widmer, G.** (2023). Automatic note-level score-to-performance alignments in the ASAP dataset. *Transactions of the International Society for Music Information Retrieval*, 6(1), 27–42.
- Rhyu, S., Kim, S., and Lee, K.** (2022). Sketching the expression: Flexible rendering of expressive piano performance with self-supervised learning. In *Proceedings of the 23rd International Society for Music Information Retrieval Conference (ISMIR)*, Bengaluru, India, pp. 178–185.
- Shi, Z., Sapp, C., Arul, K., McBride, J., and Smith III, J. O.** (2019). SUPRA: Digitizing the Stanford University Piano Roll Archive. In *Proceedings of the 20th International Society for Music Information Retrieval Conference (ISMIR)*, Delft, Netherlands, pp. 517–523.
- Simonetta, F., Avanzini, F., and Ntalampiras, S.** (2022). A perceptual measure for evaluating the resynthesis of automatic music transcriptions. *Multimedia Tools and Applications*, 81(22), 32371–32391.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y.** (2024). RoFormer: Enhanced Transformer with rotary position embedding. *Neurocomputing*, 568, 127063.
- Tang, J., Cooper, E., Wang, X., Yamagishi, J., and Fazekas, G.** (2025). Towards an integrated approach for expressive piano performance synthesis from music scores. In *Proceedings of the 50th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Hyderabad, India, pp. 1–5.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I.** (2017). Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS)* (Volume 30, pp. 5998–6008). Curran Associates, Inc.
- Watson, M.** (2018). MuseScore. *Journal of the Musical Arts in Africa*, 15(1–2), 143–147.
- Xia, G. G.** (2016). *Expressive Collaborative Music Performance via Machine Learning* [PhD thesis]. Carnegie Mellon University.
- Yan, Y., and Duan, Z.** (2024). Scoring time intervals using non-hierarchical Transformer for automatic piano transcription. In *Proceedings of the 25th International Society for Music Information Retrieval Conference (ISMIR)*, San Francisco, CA, USA, pp. 973–980.
- Ycart, A., Liu, L., Benetos, E., and Pearce, M.** (2020). Investigating the perceptual validity of evaluation metrics

for automatic piano music transcription. *Transactions of the International Society for Music Information Retrieval*, 3(1), 68–81.

- Zacharov, I., Arslanov, R., Gunin, M., Stefonishin, D., Bykov, A., Pavlov, S., Panarin, O., Maliutin, A., Rykovanov, S., and Fedorov, M.** (2019). “Zhores”: Petaflops supercomputer for data-driven modeling, machine learning and artificial intelligence installed in Skolkovo Institute of Science and Technology. *Open Engineering*, 9(1), 512–520.
- Zeng, M., Tan, X., Wang, R., Ju, Z., Qin, T., and Liu, T.-Y.** (2021). MusicBERT: Symbolic music understanding with large-scale

pre-training. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 791–800.

- Zhang, H., Chowdhury, S., Cancino-Chacón, C. E., Liang, J., Dixon, S., and Widmer, G.** (2024). DEXter: Learning and controlling performance expression with diffusion models. *Applied Sciences*, 14(15), 6543.
- Zhang, H., Tang, J., Rafee, S. R. M., and Fazekas, S. D. G.** (2022). ATEPP: A dataset of automatically transcribed expressive piano performance. In *Proceedings of the 23rd International Society for Music Information Retrieval Conference (ISMIR)*, Bengaluru, India, pp. 446–453.

TO CITE THIS ARTICLE:

Borovik, I. (2026). PianoCoRe: Combined and Refined Piano MIDI Dataset. *Transactions of the International Society for Music Information Retrieval*, 9(1), 144–163. DOI: <https://doi.org/10.5334/tismir.333>

Submitted: 17 August 2025 **Accepted:** 16 March 2026 **Published:** 27 April 2026

COPYRIGHT:

© 2026 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <https://creativecommons.org/licenses/by/4.0/>.

Transactions of the International Society for Music Information Retrieval is a peer-reviewed open access journal published by Ubiquity Press.