

JUST LABEL THE REPEATS FOR IN-THE-WILD AUDIO-TO-SCORE ALIGNMENT

Irmak Bukey
Carnegie Mellon University

Michael Feffer
Carnegie Mellon University

Chris Donahue
Carnegie Mellon University

ABSTRACT

We propose an efficient workflow for high-quality offline alignment of in-the-wild performance audio and corresponding sheet music scans (images).¹ Recent work on audio-to-score alignment extends dynamic time warping (DTW) to be theoretically able to handle *jumps* in sheet music induced by repeat signs—this method requires no human annotations, but we show that it often yields low-quality alignments. As an alternative, we propose a workflow and interface that allows users to quickly annotate jumps (by clicking on repeat signs), requiring a small amount of human supervision but yielding much higher quality alignments on average. Additionally, we refine audio and score feature representations to improve alignment quality by: (1) integrating measure detection into the score feature representation, and (2) using raw onset prediction probabilities from a music transcription model instead of piano roll. We propose an evaluation protocol for audio-to-score alignment that computes the distance between the estimated and ground truth alignment in units of measures. Under this evaluation, we find that our proposed jump annotation workflow and improved feature representations together improve alignment accuracy by 150% relative to prior work (33% \rightarrow 82%).

1. INTRODUCTION

Sheet music has been used as a primary means of communicating musical ideas for centuries. Accordingly, sheet music is a profoundly important modality for MIR, not only because of the breadth of musical knowledge and history contained within, but also because sheet music constitutes a vital interface between MIR systems and musicians. However, while multimodal MIR systems are rapidly improving at tasks like music transcription [3–6] and controllable generation [7–9], these systems typically operate on MIDI as a symbolic music format. This may be less useful to musicians, e.g., a musician might prefer transcription systems to output sheet music instead of MIDI.

¹ Video examples: <https://bit.ly/jltr-ismir2024>
Code: <https://github.com/irmakbky/jltr-alignment>
Corresponding author: Irmak Bukey <ibukey@cs.cmu.edu>

We conjecture that the scarcity of fine-grained alignment data linking sheet music to corresponding performance audio is a key bottleneck to incorporating sheet music into multimodal MIR systems. Alignments allow multimodal MIR data to be segmented into input-output chunks of tractable length for training models, and the lack of sheet music alignments may partially explain why sheet music is mostly overlooked. Moreover, alignments have practical utility outside of multimodal MIR, e.g., they may be used by musicians to practice along with pre-recorded accompaniments. Unfortunately, collecting alignments is deceptively tricky. For example, one could have a musician use a touch screen to point to the current location in sheet music while listening to a recording in real time. However, their tracking may be imprecise (due to expressive performance timing) and lack non-obvious details that are essential for segmentation (bar line locations, number of active staves).

In this work, we investigate the task of alignment of offline in-the-wild performance audio and corresponding sheet music scans (images), with a long-term goal of aligning large corpora of sheet music and performance recordings at scale. Much of the past work on audio-to-score alignment make at least one of several common assumptions that inhibit their practicality for collecting aligned data at scale: (i) the presumed availability of digital scores like MIDI or MusicXML as opposed to sheet music images [10–14], (ii) the alignment of MIDI performances or synthesized audio instead of real audio recordings [12, 15, 16], (iii) limitations in instrument diversity, commonly piano only [10, 16, 17], or (iv) dependence on time-consuming human annotation [18, 19].

Here we propose an audio-to-score alignment procedure that makes none of these assumptions, potentially offering a path forward for large-scale data collection. Most closely related to our approach is that of Shan et al. [16, 17], who examine offline alignment of in-the-wild piano sheet music images and performance recordings by aligning feature representations derived from the score and audio via MIR methods. In addition to operating on more diverse ensembles, our work has two primary distinctions: (1) we take a different approach to handling *jumps* in scores, and (2) we modify their feature representations.

A key challenge in audio-to-score alignment is handling inter-measure jumps in scores induced by repeat signs. Shan et al. [16, 17] propose extensions to DTW that are capable of automatically handling jumps. Here we propose a pragmatic alternative: a workflow and interface that allows humans to quickly annotate jumps, and a system that incor-






Figure 1. An overview of the task of audio-to-score alignment and our proposed approach. Given a score image (as a PDF) and corresponding performance audio (e.g., an MP3) as input, the task involves outputting an alignment between time in the recording and playheads in the score image. A key challenge in this task is handling jumps in the score, e.g., those created by repeat signs. In lieu of robust automatic methods for detecting or handling jumps, we propose a pragmatic approach of having experts simply label the repeats, which can be done quickly and greatly improves task performance. Our proposed system combines the repeat labels with score feature representations inspired by past work on bootleg scores [1]. This score representation is aligned with audio feature representations inspired by [2] using ordinary DTW.

porates these jump labels. We find that this approach can yield much higher-quality alignments than the automatic one, costing only seconds of annotator time.

We additionally extend the *bootleg score* feature representations used by Shan et al. [17], first proposed by Yang et al. [1]. Creating a bootleg score involves detecting noteheads and staff lines to produce a simple binary representation of a score that is conducive to alignment. We find that the use of measure bounding box detection as a preprocessing step improves the quality of underlying notehead and staff line detection algorithms. Additionally, motivated by findings in [2], we find that using raw onset probabilities predicted by a music transcription model as the audio feature representation produces higher quality alignments than using the MIDI transcriptions—see Figure 1 for a summary.

Motivated by our long-term goals of bringing sheet music into multimodal MIR, we also propose a new measure-aware evaluation scheme for comparing alignments. We speculate that measure-level alignment granularity is necessary for tractable training of multimodal MIR systems in the short term and that human perception of alignment quality is tied to measures. Accordingly, we prescribe new measure-aware alignment metrics for this task, such as an accuracy metric which reports the proportion of time where the estimated alignment is within a half measure radius of the ground truth alignment. On a small but diverse dataset of in-the-wild sheet music and aligned audio [18], we observe that our proposed system achieves an accuracy of 120% relative to that of Shan et al. (33% → 72%). By providing repeat labels, we improve the absolute accuracy of our system from 20% → 83% on a subset of pieces that have repeats. Our work makes the following contributions:

- A system capable of high-quality in-the-wild alignment of sheet music images and performance audio.
- A pragmatic workflow we call *Just Label The Repeats* that further improves alignment accuracy.
- An interface that enables rapid jump annotation.

2. TASK DESCRIPTION

Motivated by Thickstun et al. [20], here we formalize both the task of in-the-wild audio-to-score alignment and our proposed measure-aware evaluation. For a sheet music image with P pages (henceforth, a *score*), we define a *score playhead* (aligned position marker) as a tuple $(p, y, h, x) \in \mathcal{S} = \{0, \dots, P - 1\} \times [0, 1]^3$, where p is the page number, y and h are the offset and height of the current *system* (collection of staves) relative to the top edge and height of the page, and x is the playhead offset relative to the left edge of the page (see Figure 2). An analogous *audio playhead* is comparatively straightforward: a timestamp $t \in [0, T)$, where T is the audio length in seconds. An *alignment* is a mapping from audio to score playheads, i.e., $[0, T) \rightarrow \mathcal{S}$.




Figure 2. A score playhead (blue line), the output of an audio-to-score alignment, is characterized by its vertical offset (y), horizontal offset (x), and height (h), all relative to the page. A measure-aware alignment is indexed by m , a fractional measure, that can be converted to a score playhead by lookup and interpolation in a list of bounding boxes (brown outlines). Our measure-aware evaluation compares estimated playheads m' to ground truth m^* .

2.1 Measure-aware alignment

The above task definition is intended to be broad enough to encompass both past and future work on this task. Here we define a more specific form of alignment based around the location and ordering of *measures* in the score.

In this setting, measures are characterized by an ordered list of M bounding boxes $\mathbb{M} = [b_0, \dots, b_{M-1}]$, where $b_i = (b_i^p, b_i^y, b_i^h, b_i^x, b_i^w)$. Respectively, this tuple defines for each bounding box its page number, vertical offset, height, horizontal offset, and width. The ordering of this list is defined as the *logical* order that an expert would traverse when performing the piece—all *jumps* (repeat signs, Dal segno, etc.) are unrolled. For example, a score with 4 measures and a repeat implies that $M = 8$ and $b_i = b_{i+4}$.

Given a list of bounding boxes, a *measure-aware* score playhead can be characterized by a single continuous value $m \in [0, M)$, where $b_{\lfloor m \rfloor}$ is the bounding box of the current measure and fractional residual $m - \lfloor m \rfloor$ represents the offset from the left edge of the bounding box relative to its width. To convert a measure-aware score playhead in $[0, M)$ to an ordinary score playhead in \mathcal{S} , we define $h_{\mathbb{M}} : m \mapsto (b_{\lfloor m \rfloor}^p, b_{\lfloor m \rfloor}^y, b_{\lfloor m \rfloor}^h, b_{\lfloor m \rfloor}^x + b_{\lfloor m \rfloor}^w(m - \lfloor m \rfloor))$.

Given \mathbb{M} , a *measure-aware alignment* is a function $g : [0, T) \rightarrow [0, M)$. Because outputs of g index logical order (where jumps are unrolled), a measure-aware alignment g is a monotonically increasing function, i.e., $g(t_a) \leq g(t_b) \iff t_a \leq t_b$. Furthermore, we can compose $h_{\mathbb{M}}$ and g to induce an alignment that outputs score playheads, i.e., $h_{\mathbb{M}} \circ g : [0, T) \rightarrow \mathcal{S}$.

2.2 Measure-aware evaluation

Here we propose three measure-aware metrics for evaluating estimated alignments. **Our primary evaluation metric, MAcc, is defined as the proportion of time where the estimated score playhead is within a half measure radius of the ground truth score playhead**, which we posit is sufficiently precise for broader goals of multimodal MIR systems. ME_{Err} and MD_{Dev} are the mean and standard deviation (across time) of the absolute error between the estimated and ground truth playheads in units of ground truth measures.

More formally, given a ground truth measure-aware alignment g^* characterized by measures \mathbb{M}^* , and an estimated alignment g' characterized by \mathbb{M}' , we define:

$$\begin{aligned} \text{MDiff}(t) &= \text{Reindex}(g'(t), \mathbb{M}', \mathbb{M}^*) - g^*(t), \\ \text{MAcc} &\equiv \frac{1}{N} \sum_{i=0}^{N-1} \begin{cases} 1 & \text{if } |\text{MDiff}(\frac{Ti}{N})| \leq \frac{1}{2} \\ 0 & \text{otherwise.} \end{cases} \\ \text{MErr} &\equiv \frac{1}{N} \sum_{i=0}^{N-1} \left| \text{MDiff}\left(\frac{Ti}{N}\right) \right|, \\ \text{MDev} &\equiv \sqrt{\frac{1}{N} \sum_{i=0}^{N-1} \text{MDiff}\left(\frac{Ti}{N}\right)^2} \end{aligned}$$

We set $N = 100T$, i.e., we compute all metrics at a resolution of 100 comparisons per second. As an example, Figure 2 shows a single comparison where $\text{MDiff} = 0.25$.

The `Reindex` procedure ensures that the evaluation is based in units of ground truth measure indices, despite potential discrepancies between \mathbb{M}' and \mathbb{M}^* . Informally, this procedure matches each box in \mathbb{M}' to the box in \mathbb{M}^* that it is closest to in terms of Euclidean distance between mid-points, using the ordering from \mathbb{M}^* to break ties that occur for repeated measures.

3. SYSTEM DESCRIPTION

In this section, we detail our proposed system for in-the-wild audio-to-score alignment. At a high level, our method uses DTW to align piano roll-like feature representations extracted independently from the audio and score (Figure 1). Both representations are matrices where one axis corresponds to time (either in units of seconds or measures) and the other corresponds to 88 MIDI pitches from A0 to C8 (piano range). Specifically, score feature representations are matrices in $\{0, 1\}^{48M \times 88}$ (binary) where M is the number of measures in the score, and audio feature representations are matrices in $[0, 1]^{Tf_k \times 88}$ (continuous) for performance length T and frame rate $f_k = 31\text{Hz}$.

3.1 Score feature representation

Our proposed score feature representation is an extension of *bootleg scores* proposed by Yang et al. [1]. Extracting bootleg scores involves detecting noteheads and staff lines, and then combining this information into a binary matrix where a 1 encodes the presence of a notehead at a particular horizontal position on a particular staff line. Our system extends this representation in two ways: (1) we use measure detection as a preprocessing step and run notehead detection algorithms on segmented measure images instead of full pages, and (2) we translate notehead positions on staff lines into MIDI pitches before alignment.

Measure, notehead and staff line detection. We use the methods from [1] to detect noteheads and staff lines from score images. Instead of operating on full page images, we first segment pages into measures using the measure detection model from [21], and detect noteheads and staff lines on individual measure images resized to fixed dimensions. In preliminary experiments, we found that measure segmentation improved detection consistency—our intuition is that these methods are sensitive to absolute pixel sizes of noteheads and stafflines, and resizing individual measures to uniform size reduces variance in the sizes of these attributes across measures and pieces. Here we resize measures by resizing the smaller of their height and width to 900 pixels, preserving aspect ratio for the larger of the two. For each notehead, we retain its *bootleg location*, defined as its raw (pixel-wise) horizontal location within the measure, and its semantic (discrete) vertical position within the detected staff lines (e.g., an F and G natural in treble clef are one apart in their *staff positions*).

Piano rolls. We create a binary piano roll-like representation of the score using the bootleg locations of noteheads. Specifically, for each logical measure image index $k \in \{0, \dots, M-1\}$, we construct a binary matrix $S_k \in \{0, 1\}^{48 \times 88}$, where a 1 at row i and column j corresponds to a notehead with horizontal location $\frac{j}{48}$ relative to the

measure width, and its staff position converted to a pitch j . We pick a measure representation to have 48 rows to give sufficient resolution to a variety of rhythmic patterns and note durations. We concatenate $[S_0, \dots, S_{M-1}]$ together to form our final representation $S \in \{0, 1\}^{48M \times 88}$.

Converting staff position to pitch. A key obstacle is that, without key signature and clef information, the mapping from staff positions to MIDI pitches is ambiguous. If we had these attributes, we could simply look up the pitch associated with each staff line. However, we found existing OMR systems to have brittle support for detecting this information for in-the-wild sheet music—hence, our method does not assume that we have access to this information. Accordingly, we convert bootleg scores into piano rolls by simply assuming treble and bass clefs respectively when two staves are detected—if more or fewer staves are detected, we default to the treble clef—and the key of C major. Surprisingly, perhaps because of the global optimality of DTW, these assumptions lead to reasonable alignments even when they are incorrect. We note that ground truth key signature and clef information for each measure can be fed into our method to facilitate and improve this conversion, but we do not require it.

3.2 Audio feature representation

Our audio feature representation pipeline is comparatively simple. To compute it, we simply pass the audio through the Onsets and Frames piano transcription model [3]. Motivated by [2], we use the raw onset prediction probabilities from this model as our audio feature representation, which is a matrix in $[0, 1]^{Tf_k \times 88}$. Despite this transcription model being trained on piano, we find that its onset probabilities can yield reasonable alignments even for non-piano audio.

3.3 Alignment

Finally, we align the score representations in $\{0, 1\}^{48M \times 88}$ and audio representations in $\mathbb{R}^{Tf_k \times 88}$. We use the implementation of standard DTW from `librosa` [22] with default parameters: equal-weighted transitions $(1, 1)$, $(0, 1)$, and $(1, 0)$, and Euclidean distance to compute costs.

4. EXPERIMENTS

Here we detail our experiments, which center around comparing our proposed method to that of Shan et al. [17] on the MeSA-13 [18] and SMR [1] datasets using our proposed measure-level evaluation (see Section 2.2).

4.1 Datasets

We evaluate our approach and relevant baselines on two different datasets. The first is MeSA-13 [18], a dataset of 13 sheet music scans and corresponding real (i.e., not synthesized) performance audio. This dataset contains expert annotations of measure bounding boxes in logical order (\mathbb{M}^*) and the timestamp of every measure in the performance audio (we linearly interpolate between timestamps to get a continuous ground truth mapping g^*). While small, MeSA-13 has reasonable diversity in score typesetting, performance acoustics (two pieces feature instruments besides piano), and jumps (two pieces have repeats).

The second dataset is a subset of the Sheet MIDI Retrieval v1.0 (SMR) dataset [1]. The full dataset contains scanned scores from IMSLP for 100 solo piano pieces (none of which have jumps), corresponding MIDI performances synthesized as audio, and human annotations of measures per line and measure timestamps. Of notable absence are annotations of measure bounding boxes, which are required for our proposed evaluation. Accordingly, we detect measures [21] and discard pieces where the detections do not agree with annotations of measures per line—this leaves us with a subset of 60 pieces for evaluation. Henceforth, SMR refers to this subset.

4.2 Access to additional annotations

We primarily evaluate systems in an automatic setting where systems are only given the score and audio as input. Because our system can incorporate additional score annotations when available, we also evaluate in settings where our system has access to additional annotations from the ground truth, simulating workflows where experts are in the loop during alignment. Specifically, we explore settings where our method has access to ground truth repeat annotations (R), measure bounding boxes (M), and staff information (S)—clef and key signatures. We only evaluate in these settings on MeSA-13 where we have these labels.

4.3 Baselines

We compare the performance of our system, composed of our feature extraction pipeline and vanilla DTW, to that of [17]. In the latter system, the feature extraction pipeline uses bootleg scores and staff line detection on the score images to extract staff lines (referred to as segments). For audio features, a transcribed MIDI representation is obtained from the Onsets and Frames piano transcription model [3] which is then used to compute bootleg scores. Finally, Hierarchical DTW performs a segment-level alignment between score features and audio features while handling jumps and repeats that occur at segment boundaries, but not those within segments. Thus, we opt to evaluate this baseline approach at the measure level instead of the segment level (which also allows for comparison with our system) by converting the segment-level alignment to a measure-level one via an algorithm with several key steps.

We first use measure detection [21] to locate measures in each segment. Then, we map segment indices to measure indices based on the positions of detected measures. Finally, we turn the given alignment between audio timestamps and segment indices to one between audio timestamps and measure indices via linear interpolation.

We compare these two systems as proposed instead of comparing their components for two reasons. First, in [17] it is claimed the system performs segment-level alignment on pieces without repeat info; we aim to test this. Second, while Hierarchical DTW allows backwards jumps to prior segments, it only allows forward jumps to one segment past the last one seen. Using Hierarchical DTW at the measure level would limit possible forward jumps to only one measure past the last one observed, which is insufficient for realistic alignment tasks.

Dataset	Given	System	MAcc	MErr	MDev
M13	-	[17]	0.33	10.9	11.6
	-	Ours	0.72	1.9	3.7
	R [†]	Ours	0.82	0.4	0.2
	R,M	Ours	0.86	0.4	0.2
	R,M,S	Ours	0.88	0.3	0.2
M13 _R	-	[17]	0.17	23.6	10.2
	-	Ours	0.20	10.0	3.0
	R [†]	Ours	0.83	0.3	0.0
	R,M	Ours	0.93	0.2	0.0
	R,M,S	Ours	0.95	0.2	0.0
M13 _{NR}	-	[17]	0.36	8.6	10.2
	-	Ours	0.82	0.4	0.2
	R [†]	Ours	0.82	0.4	0.2
	R,M	Ours	0.85	0.4	0.3
	R,M,S	Ours	0.87	0.3	0.2
SMR	-	[17]	0.36	14.2	18.7
	-	Ours	0.82	2.9	18.8

Table 1. Evaluation on MeSA-13 (including subsets with Repeats and No Repeats) and SMR. Our method outperforms that of [17] across all datasets except the subset of MeSA-13 with no repeats. R[†] is our recommended setting where our method is given access to ground truth Repeats that require little time for humans to annotate—we observe limited gains from more time-consuming annotations of Measure bounding boxes and Staff metadata.

4.4 Results and discussion

In Table 1, we report the measure-level alignment metrics (Section 2.2) of our system in all four settings and the system of [17] across the MeSA-13 (M13) and SMR datasets. To emphasize the effect of jumps on alignment performance, we separately report performance on the subset of MeSA-13 pieces with and without repeats (M13_R and M13_{NR}, respectively).

We observe that our system outperforms that of [17] in the automatic setting across all datasets. The superior performance of our system over that of [17] is likely due to our refinements to feature representations. We also note here that the system of [17] is designed to work on line-level, therefore evaluating it using our measure-level metric yields lower accuracy than what was reported in [17].

Additionally, we observe that given repeats and ground truth measure annotations, our system’s performance improves by 22% relative (MAcc 0.72 → 0.88) on M13. However, we also observe a relative performance improvement of 14% (MAcc 0.72 → 0.82) using our system on the same dataset when we only pass in repeats. Given that repeats are much easier for humans to annotate than measure bounding boxes, we propose to have humans *just label the repeats* as a recommended tradeoff between alignment quality and annotator time. We also explore providing our system with measure-level key signature and clef

Representation	M13	SMR
Onset probabilities	0.88	0.82
Onset predictions	0.86	0.82
Frame probabilities	0.70	0.53
Frame predictions	0.66	0.51
MIDI	0.46	0.20

Table 2. Evaluation of measure-aware alignment accuracies (MAcc) achieved by different audio feature representations obtained from the Onsets and Frames piano transcription model [3] on MeSA-13 and SMR.

information, finding that this information only marginally improves performance relative to the default key and clef assumptions described in Section 3.1.

4.5 Different audio feature representations

Here we compare alternative audio feature representations by evaluating MAcc on M13 given all additional information (i.e., the R,M,S setting described in Section 4.2). While we primarily use raw onset prediction probabilities (Section 3.2), the Onsets and Frames model [3] provides other possibilities including onset predictions (thresholded probabilities), frame probabilities and predictions, and the postprocessed MIDI transcription converted to piano roll (see [3] for details). Table 2 shows that onsets consistently outperform frames as an alignment representation—our intuition is that onsets are more appropriate for our setting as the bootleg score representation does not encode note durations. Additionally, while transcribed MIDI is a common feature representation for music alignment, in our setting we find it to be the worst choice.

5. LABELING INTERFACE

Here we describe a web-based interface that we built to enable experts to quickly annotate jumps in scores (induced by repeat signs, D.S. al coda, etc.). Our experiments show large improvements in alignment quality given jump labels. Accordingly, we designed an interface to make this process efficient—experts can label jumps in a matter of seconds. We include videos demonstrating the end-to-end process and qualitative results of our proposed workflow for pieces outside our evaluation data.² In these videos, labeling jumps takes less than 6s per page on average.

Our interface features a unified workflow for jump annotation based on clicking the starting and ending measure of a jump (Figure 3). To enable this workflow, we first run measure detection [21] on the backend and visualize detected measures on the frontend as a user hovers over the score. Then, users can simply click on two different measures to set a jump—this simple unified workflow accommodates a long tail of jump glyphs. The interface also visualizes the logical order of the measures induced by the measure bounding boxes and any jumps the user has set (M from Section 2.1). Finally, a user can download the

² Video examples: <https://bit.ly/jltr-ismir2024>

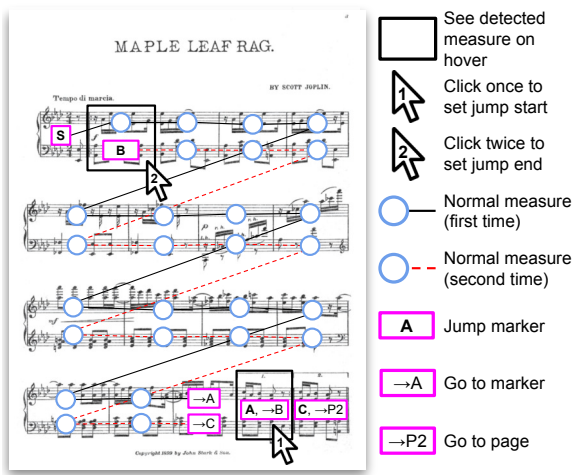


Figure 3. Illustration of our web-based interface for labeling jumps (e.g. repeats) in scores. Our interface enables rapid jump annotation (just seconds per page after training), which we find to dramatically improve alignment quality on pieces with jumps.

logical-order measures as a simple JSON file, which can then be loaded into our Python-based alignment package.

6. RELATED WORK

Our work relates to prior work in offline audio-to-score alignment, score following, and score annotation tools.

6.1 Offline alignment

Offline alignment entails building a correspondence between sheet music and performance recording of the contained music. As described previously, our work is inspired by the bootleg score line of research [1, 16, 17, 23, 24] as well as the usage of raw audio prediction scores in [2]. Our contributions stem from combining these approaches with minimal human annotations to outperform prior baselines. Other prior work in this area includes [15, 25], which use MIDI events as an intermediate representation between scores and audio, and [26–31], which use chroma vectors (where components correspond to the 12 pitch types in Western music). We perform a similar procedure to the methods used in these works except that we use bootleg scores as our intermediate representation. We diverge from [14], which uses MusicXML as an intermediate representation, and from [32], which uses LilyPond representations, but we mention them here as related approaches. We also incorporate reasoning from [20] regarding evaluation metrics.

6.2 Score following

In contrast to offline alignment, *score following* involves building a real-time alignment between sheet music and live performance audio. Initial solutions to this problem include [33, 34], with later works addressing jumps and repeats [35, 36]; for more on related work in past decades, see survey papers [37, 38]. Unlike most other research in score following which often assumes that a digital score

representation (like MIDI) is available, our emphasis is on *solely using score images* with an aim to perform alignment at scale. This said, some recent work does attempt to solve this problem in images. For instance, works such as [39–42] map audio snippets to corresponding places in score images using neural networks, but they are limited to piano music. We diverge from them by considering a range of different types of raw score images and audio (such as ones with instrumentation beyond solo piano) and leveraging bootleg scores [1, 16, 17, 23, 24] for mapping, but these are still related to our work.

6.3 Sheet music annotation interfaces

Our work also relates to past work on designing interfaces to assist in the annotation of sheet music. Most directly related is that of Feffer et al. [18] which attempts to facilitate interactive annotation of sheet music and audio alignment via a workflow based on aligning detected beat timestamps [43] to detected measures [21]. This interface was used to compile the MeSA-13 dataset of aligned audio and scores, which we use to evaluate our work, though we note that their interface required 20 hours of expert time to collect less than an hour of aligned data. In contrast, our interface is designed to be used for a matter of seconds to annotate repeats. Other interfaces focus on facilitating measure bounding box annotations [44, 45], which is complementary to our workflow that focuses on repeat annotation using predicted bounding boxes. Lastly, Soundslice [19] is a commercial product that offers an interactive alignment workflow based on stronger notions of OMR, but its implementation details are proprietary.

7. CONCLUSION

In summary, we introduce a workflow for efficiently aligning sheet music images to performance audio. The key insight we leverage is that while automated alignment algorithms are currently not robust to repeats in scores, humans can quickly label these repeats, thereby improving alignment performance. We validate this approach on a dataset of in-the-wild sheet music scans and real performance recordings, showing that we outperform existing baselines that only use automated approaches.

Given these results, one future project we aim to undertake is to collect jump annotations at scale to create large aligned datasets. We acknowledge that the datasets we used to evaluate our approach are small, but the insights gained from them can help scale up data for future evaluations. For instance, we could extend the interface from Section 5 to allow annotators to quickly audit and adjust alignments to collect more data, as in [18]. Moreover, other future work could revisit the creation of a fully automated alignment algorithm with insights from our work, namely that such an algorithm that leverages OMR to identify jumps and repeats may be more successful than one that does not. Collecting more data would therefore be helpful for developing and evaluating future approaches. Lastly, as described in the start of our paper, large aligned datasets could be used to derive multimodal MIR systems for music students and professionals alike.

8. ETHICS STATEMENT

As described previously, advancements in audio-to-score alignment can result in new multimodal datasets derived from existing repositories of sheet music and audio (such as IMSLP [46]) and new interactive music systems tailored for performance. Our motivation for pursuing this direction is to unlock multimodal MIR systems that (1) supplement music education by helping performers rehearse, (2) understand or generate sheet music to unlock seamless communication with human musicians, and (3) reduce reliance on copyrighted material for building music AI (i.e., by leveraging public domain scores and recordings).

We also recognize several potential ethical concerns stemming from our work. Firstly, our method is firmly rooted in conventions of Western music. Accordingly, downstream systems and data derived from our method may reflect a Western bias that does not generalize well to other musical traditions, especially those with different notation or tuning systems. Secondly, though our goal is to lessen the amounts of copyright infringement taking place to build generative music AI, multimodal MIR systems could be used to circumvent data protections, e.g., by transcribing copyrighted recordings as less-protected sheet music. Lastly, increased ability to understand sheet music could lead to deepfakes or misinformation, e.g., scores that could be falsely attributed to Beethoven, or ragtime recordings that could be falsely attributed to Joplin. In response to these concerns, we recommend that future work mitigate these risks by, for example, developing analogous systems capable of improving understanding of non-Western music notation. We also recommend that MIR researchers should be mindful of data protections, copyright violations, and artistic mimicry that, if subverted, could threaten the livelihood of musicians.

9. REFERENCES

- [1] D. Yang, T. Tanprasert, T. Jenrungrot, M. Shan, and T. J. Tsai, "Midi passage retrieval using cell phone pictures of sheet music," in *ISMIR*, 2019.
- [2] B. Maman and A. H. Bermano, "Unaligned supervision for automatic music transcription in the wild," in *International Conference on Machine Learning (ICML)*, 2022.
- [3] C. Hawthorne, E. Elsen, J. Song, A. Roberts, I. Simon, C. Raffel, J. Engel, S. Oore, and D. Eck, "Onsets and frames: Dual-objective piano transcription," *ISMIR*, 2018.
- [4] C. Donahue, J. Thickstun, and P. Liang, "Melody transcription via generative pre-training," in *ISMIR*, 2022.
- [5] J. Gardner, I. Simon, E. Manilow, C. Hawthorne, and J. Engel, "Mt3: Multi-task multitrack music transcription," in *International Conference on Learning Representations (ICLR)*, 2022.
- [6] R. M. Bittner, J. J. Bosch, D. Rubinstein, G. Meseguer-Brocal, and S. Ewert, "A lightweight instrument-agnostic model for polyphonic note transcription and multipitch estimation," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2022.
- [7] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzetti, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi *et al.*, "MusicLM: Generating music from text," *arXiv:2301.11325*, 2023.
- [8] J. Copet, F. Kreuk, I. Gat, T. Remez, D. Kant, G. Synnaeve, Y. Adi, and A. Défossez, "Simple and controllable music generation," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [9] S.-L. Wu, C. Donahue, S. Watanabe, and N. J. Bryan, "Music controlnet: Multiple time-varying controls for music generation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [10] J. J. Carabias-Orti, F. J. Rodriguez-Serrano, P. Vera-Candeas, N. Ruiz-Reyes, and F. J. Canadas-Quesada, "An audio to score alignment framework using spectral factorization and dynamic time warping," in *ISMIR*, 2015.
- [11] J.-L. Syue, L. Su, Y.-J. Lin, P.-C. Li, Y.-K. Lu, Y.-L. Wang, A. W. Su *et al.*, "Accurate audio-to-score alignment for expressive violin recordings," in *ISMIR*, 2017.
- [12] A. Arzt and S. Lattner, "Audio-to-score alignment using transposition-invariant features," in *ISMIR*, 2018.
- [13] T. Tanprasert, T. Jenrungrot, M. Müller, and T. Tsai, "Midi-sheet music alignment using bootleg score synthesis," in *ISMIR*, 2019.
- [14] F. Foscarin, A. Mcleod, P. Rigaux, F. Jacquemard, and M. Sakai, "Asap: a dataset of aligned scores and performances for piano transcription," in *ISMIR*, 2020.
- [15] M. Dorfer, A. Arzt, and G. Widmer, "Learning audio-sheet music correspondences for score identification and offline alignment," in *ISMIR*, 2017.
- [16] M. Shan and T. Tsai, "Improved handling of repeats and jumps in audio-sheet image synchronization," in *ISMIR*, 2020.
- [17] ———, "Automatic generation of piano score following videos." *Transactions of the International Society for Music Information Retrieval (TISMIR)*, 2021.
- [18] M. Feffer, C. Donahue, and Z. Lipton, "Assistive alignment of in-the-wild sheet music and performances," in *ISMIR Late Breaking Demo Track*, 2022.
- [19] Soundslice LLC, "Soundslice: Create living sheet music," <https://www.soundslice.com>, 2024, accessed on 2024-04-12.

- [20] J. Thickstun, J. Brennan, and H. Verma, “Rethinking evaluation methodology for audio-to-score alignment,” *arXiv preprint arXiv:2009.14374*, 2020.
- [21] S. Waloschek, A. Hadjakos, and A. Pacha, “Identification and cross-document alignment of measures in music score images,” in *ISMIR*, 2019.
- [22] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, “librosa: Audio and music signal analysis in Python,” in *SciPy*, 2015.
- [23] T. Tsai, D. Yang, M. Shan, T. Tanprasert, and T. Jen-rungrot, “Using cell phone pictures of sheet music to retrieve midi passages,” *IEEE Transactions on Multimedia*, 2020.
- [24] D. Yang, A. Goutam, K. Ji, and T. Tsai, “Large-scale multimodal piano music identification using marketplace fingerprinting,” *Algorithms*, 2022.
- [25] M. Dorfer, J. Hajic Jr, A. Arzt, H. Frostel, and G. Widmer, “Learning audio-sheet music correspondences for cross-modal retrieval and piece identification.” *Transactions of the International Society for Music Information Retrieval (TISMIR)*, 2018.
- [26] F. Kurth, M. Müller, C. Fremerey, Y.-h. Chang, and M. Clausen, “Automated synchronization of scanned sheet music with audio recordings,” in *ISMIR*, 2007.
- [27] D. Damm, C. Fremerey, F. Kurth, M. Müller, and M. Clausen, “Multimodal presentation and browsing of music,” in *Proceedings of the 10th international conference on Multimodal interfaces*, 2008.
- [28] C. Fremerey, M. Clausen, S. Ewert, and M. Müller, “Sheet music-audio identification.” in *ISMIR*, 2009.
- [29] V. Thomas, C. Fremerey, M. Müller, and M. Clausen, “Linking sheet music and audio-challenges and new approaches,” *Multimodal Music Processing*, 2012.
- [30] C. Fremerey, M. Müller, and M. Clausen, “Handling repeats and jumps in score-performance synchronization,” in *ISMIR*, 2010.
- [31] M. G. M. Gasser, A. Arzt, and G. Widmer, “Automatic alignment of music performances with structural differences,” in *ISMIR*, 2013.
- [32] L. Liu, V. Morfi, and E. Benetos, “Joint multi-pitch detection and score transcription for polyphonic piano music,” in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [33] R. B. Dannenberg, “An on-line algorithm for real-time accompaniment,” in *Proceedings of the International Computer Music Conference (ICMC)*, 1984.
- [34] B. Vercoe, “The synthetic performer in the context of live performance,” in *Proceedings of the International Computer Music Conference (ICMC)*, 1984.
- [35] B. Pardo and W. Birmingham, “Modeling form for on-line following of musical performances,” in *Proceedings of the National Conference on Artificial Intelligence*, 2005.
- [36] A. Arzt, G. Widmer, and S. Dixon, “Automatic page turning for musicians via real-time machine listening,” in *ECAI 2008*. IOS Press, 2008.
- [37] M. Puckette and C. Lippe, “Score following in practice,” in *Proceedings of the International Computer Music Conference (ICMC)*, 1992.
- [38] N. Orio, S. Lemouton, and D. Schwarz, “Score following: State of the art and new developments,” in *Proceedings of the New Interfaces for Musical Expression (NIME) Conference*, 2003.
- [39] M. Dorfer, A. Arzt, and G. Widmer, “Towards score following in sheet music images,” in *ISMIR*, 2016.
- [40] F. Henkel and G. Widmer, “Real-time music following in score sheet images via multi-resolution prediction,” *Frontiers in Computer Science*, 2021.
- [41] F. Henkel, R. Kelz, and G. Widmer, “Learning to read and follow music in complete score sheet images,” in *ISMIR*, 2020.
- [42] ———, “Audio-conditioned u-net for position estimation in full sheet images,” in *2nd International Workshop on Reading Music Systems*, 2019.
- [43] S. Böck, F. Korzeniowski, J. Schlüter, F. Krebs, and G. Widmer, “madmom: a new Python Audio and Music Signal Processing Library,” in *Proceedings of the 24th ACM International Conference on Multimedia*, 2016.
- [44] F. Zalkow, A. V. Corrales, T. Tsai, V. Arifi-Müller, and M. Müller, “Tools for semi-automatic bounding box annotation of musical measures in sheet music,” in *ISMIR Late Breaking Demo Track*, 2019.
- [45] E. Egozy and I. Clester, “Computer-assisted measure detection in a music score-following application,” in *Proceedings of the 4th International Workshop on Reading Music Systems, Online*, 2022.
- [46] “Imslp: Petrucci music library.” [Online]. Available: <https://imslp.org/>