

Improving Real-time Score Following in Opera by Combining Music with Lyrics Tracking

Charles Brazier¹ Gerhard Widmer^{1,2}

¹Institute of Computational Perception, Johannes Kepler University Linz, Austria

²LIT AI Lab, Linz Institute of Technology, Austria

firstname.lastname@jku.at

Abstract

Fully automatic opera tracking is challenging because of the acoustic complexity of the genre, combining musical and linguistic information (singing, speech) in complex ways. In this paper, we propose a new pipeline for complete opera tracking. The pipeline is based on two trackers. A music tracker that has proven to be effective at tracking orchestral parts, will lead the tracking process. In addition, a lyrics tracker, that has recently been shown to reliably track the lyrics of opera songs, will correct the music tracker when tracking parts that have a text dominance over the music. We will demonstrate the efficiency of this method on the opera *Don Giovanni*, showing that this technique helps improving accuracy and robustness of a complete opera tracker.

1 Introduction and Contribution

Score following aims at aligning classical music performances with their corresponding scores (sheet music), in order to assign a score position at each time step in the performance. There has been constant progress in this domain, starting with the tracking of monophonic melodies in (Dannenberg, 1984), all the way to recent systems that can follow, under real conditions, complex orchestral works (Arzt and Widmer, 2015) in a completely autonomous process. This has led to the development of new applications such as automatic page-turning for pianists (Arzt et al., 2008), live performance visualization (Lartillot et al., 2020), or score viewing and automatic contextualization in orchestra concerts (Prockup et al., 2013; Arzt et al., 2015) to enrich the viewers' experience.

Tracking live opera performances would become an essential tool for all future opera halls, supporting functionalities like fully automatic subtitles display, or automatic camera control and video editing for live streaming services. How-

ever, and compared to previous existing works, operas are more challenging to track, due to the setup with a complete orchestra and singers that act and sing on stage, one or several at a time, for several hours, with various noises, acting breaks, intermittent applause, musical (sometimes improvised) interludes, etc.

First attempts at opera tracking (Brazier and Widmer, 2020b,a) use an *On-Line Dynamic Time Warping (OLTW)* algorithm (Dixon, 2005) to align complete performances with a *reference performance* (some other recording of the work in question) that has been aligned to the score beforehand and serves as a proxy to the score. This audio-to-audio alignment strategy is an elegant way to circumvent the unavailability of complete opera score files in symbolic format. Also, using a real recording is advantageous because the sounds in the reference are much more realistic and similar to what is to be expected in the real performance than anything one could synthesize from a score. Brazier and Widmer (2020b) combine alignment with three audio event detectors for music, speech/singing voice, and applause, which halt the tracking process during long silences, applause, or interlude passages that can occur in between the parts. Brazier and Widmer (2020a) further improve tracking accuracy by using two trackers working in parallel, one using audio features tuned on orchestral music (Gadermaier and Widmer, 2019), the other using features tuned on the *recitative* subset of one opera performance.

In this work, we propose to exploit an additional source of information: the *lyrics* sung or spoken in the audio recordings. We do not assume the written lyrics to be available in textual form. Rather, the idea is to train an acoustic phoneme recognition model that extracts phoneme sequence estimates both from the reference (off-line) and the live performance (on-line), and to align these in

Conductor	Place	Year	Duration	Role
H.v. Karajan	Berlin	1985	2:57:53	Reference
Á. Fischer	Vienna	2018	3:12:54	Target
A. Manacorda	Vienna	2019	3:07:09	Target

Table 1: Dataset used in this study.

real time, giving us a real-time lyrics tracking algorithm. More specifically, the acoustic model will predict, for each audio frame, a probability vector over a set of phonemes. For each part (aria, recitative, etc.) in the score, we assign a *voice on music ratio* value, calculated on the reference recording with the help of dedicated music/speech audio classifiers. The music tracker leads the alignment process. As soon as the score position corresponds to a voice-dominant part in the score, the lyrics tracker starts and we rely on its score position. When the score position reaches a music-dominant part, the lyrics tracker is stopped and the music tracker alone is used.

Acoustic model and lyrics tracker have already been presented in a recent publication (Brazier and Widmer, 2021b), but only evaluated on selected text-heavy *recitativo* passages. Here we demonstrate, for the first time, the benefit of combining lyrics with music tracking in an automated way.

2 Data Description

Score followers are evaluated by computing their alignment accuracy on audio performances that have been manually annotated to the corresponding score. As no such dataset exists for opera, we had to create our own. The dataset focuses on the opera *Don Giovanni* by W.A.Mozart. As the *reference*, serving as a proxy to the score, we selected a commercial CD recording conducted by Herbert von Karajan in 1985. As *target performances* that we want to align to the score in real time, we use two full live performances, with different casts and stagings, that have been recently recorded at and by the *Vienna State Opera*, one conducted by Adam Fischer in 2018, the other by Antonello Manacorda in 2019. There are two parts in the reference that are not played in the two live recordings. For this study, we decided to remove these to align performances that follow the same score structure. Compared to the reference, the live performances contain applause, breaks, and interludes that can appear between parts. The dataset details are given in Table 1.

For each performance in the dataset, we manually affixed 5,304 bar annotations, 2,866 for the first act and 2,438 for the second, corresponding to the total number of bars present in the 500-pages score book. The annotations in the reference performance permit to link the complete performance to the score book. The annotations in the target performances serve for evaluating the alignment accuracy of our tracker. Thus, our dataset comprises more than 9 hours of opera recordings played and sung in real conditions by different orchestras and singers and recorded with different recording setups. It contains around 16,000 manual bar-level annotations assigned to the 530 pages score book, which is available online thanks to the *Mozarteum Foundation Salzburg*¹. Precisely, annotating these 9 hours of music took about 300 hours of work.

3 Real-Time Opera Trackers

Operas are complex works that combine music, singing, and speech in complex ways. Most of the time, the piece is led by the music, with singers singing on top of the orchestra. However, operas also include passages, such as *recitativo* sections, where the dominant signal is the lyrics spoken or sung by the singers, with a sparse musical accompaniment that is played differently across performances (e.g., arpeggiated chords, not aligned to the lyrics, partly improvised, and played by different instruments). To tackle this, we propose to alternate between two trackers, one focusing on the music information and the other on the lyrics information. We first describe our *music tracker* that serves as a baseline in this study and that leads the tracking process. We then describe our *lyrics tracker*, and then propose one simple way of combining them for achieving a better global tracking accuracy. This combination strategy will be experimentally verified in the next chapter.

3.1 Music Tracker

The *music tracker* is based on an adaptive version of the On-Line Time Warping (OLTW) algorithm (Dixon, 2005) that has been successfully used in orchestra (Arzt et al., 2008) and also in opera tracking (Brazier and Widmer, 2020b,a). The OLTW algorithm updates an accumulated cost vector that has the length of the reference feature sequence, where the index of its minimal

¹<https://dme.mozarteum.at/DME/nma/>

value corresponds to the score position given by the algorithm. Per audio frame, it receives as input a feature vector of 100 MFCCs (120 MFCCs are calculated from the audio sampled at 44.1 kHz, but the first 20 are discarded (Gadermaier and Widmer, 2019)), computed with a window size of 20 ms, and a hop size of 10 ms. The features of the reference audio are computed beforehand, while those of the target performance are computed in real-time. For each new incoming target feature, we compute the cosine distance between the feature and an interval of reference features of length c , centered around the expected score position (in practice c is fixed to 4000, corresponding to a context of 40 seconds of audio). Then, considering the previous score position sp , the previously accumulated cost vector D_{prev} , and the current distance vector d , we compute the value of the new accumulated cost vector D by first initializing its values by $+\infty$, and then applying the following recursive formula:

$$\forall i \in [sp - c/2 : sp_{j-1} + c/2],$$

$$D[i] = d[i - (sp - c/2)] + \min \begin{cases} D_{prev}[i - 1] \\ D_{prev}[i] \\ D[i - 1] \end{cases} \quad (1)$$

To compare costs in D among themselves and not favor shorter paths over longer ones, we normalize them by dividing all values by their distance from the initial score position (i.e. by the sum of their index in the accumulated vector and an incremental counter representing the number of iterations since the beginning of the tracking).

Our target performances are performed under real conditions and thus include applause, breaks, or interludes that can be played in between the parts. We make use of the applause, music, and speech detectors detailed in (Brazier and Widmer, 2020a) to halt the tracking process when detected.

3.2 Lyrics Tracker

The *lyrics tracker* makes use of an on-line audio-to-lyrics alignment method that has been shown to robustly track the lyrics of different languages, in the genre of opera (Brazier and Widmer, 2021b). The tracker is composed of an acoustic model that generates, in real-time, *posteriograms* representing the frame-wise probability distribution over a set of predefined phonemes through time. Then, it employs the same OLTW algorithm described

in Section 3.1, but in this case, aligning the posterio-gram of the reference performance generated beforehand, and the posterio-gram of the target performance generated online. This obviates the need for a text-to-phoneme tool to translate the written-out lyrics, as well as a manual alignment of the lyrics to the reference performance. It works without having the lyrics themselves and can track a language other than the language(s) the acoustic model was trained on, as shown in (Brazier and Widmer, 2021b).

The *acoustic model* is the core element of our lyrics tracker; its role is to estimate in real-time a posterio-gram matrix from the audio recording. Its architecture is the CP-ResNet (Koutini et al., 2019), composed of convolutional layers with residual connections between layers, and has a receptive field of 57 frames in the input feature sequence centered around its time position, fixing the delay of the model to 28 frames. The model takes as input 80 MFCCs that are extracted from an audio window of 20 ms, sampled at 16 kHz, with a hop size of 10 ms; it outputs a vector every 40 ms. The output vector is of length 60, representing the classes of the 57 different phonemes that are included in the multilingual DALI dataset (Meseguer-Brocal et al., 2018) used to train the model. The dataset collects 275 hours of Western musical genres with lyrics annotations at the sentence, word, or note level, and includes English, German, French, Spanish and Italian languages. The phoneme representation permits to train a single model on different languages (Vaglio et al., 2020). The output vector also adds the space token, the instrumental token, and the blank token, essential to a Connectionist Temporal Classification (CTC) training (Graves et al., 2006) (the blank class will be ignored when applying Equation 1).

3.3 M&L Tracker: Combining Music and Lyrics Trackers

To exploit the complementarity between the two previously described trackers, we first classify each part of the opera in two classes (in the given reference performance): parts dominated by the music and parts dominated by the voice. To do so, we use the structure detailed in the Table of Contents of the Opera², and consider each title as an individual part. For each part, we use the mu-

²dme.mozarteum.at/DME/nma/nma_toc.php?vsep=68

sic and voice detectors (already used to halt the tracking process in between parts, as mentioned in Section 3.1 above) to calculate a *voice over music ratio* that is given by the percentage of voice along the part divided by the percentage of music. Thus, an instrumental part will have a ratio close to 0, whereas a part that contains more voice than music will have a ratio higher than 1.

The combination of the two proposed tracking models is delicate because they both work at a different pace (10ms for the music tracker, and 40ms for the lyrics tracker), the lyrics tracker has a delay of 280ms in its output due to its receptive field, and neither of them is able to track accurately full opera performances. More precisely, the music tracker is inaccurate when an improvised accompaniment is played during a part led by the lyrics, and the lyrics tracker is entirely lost during instrumental parts. Our approach is to use the music tracker continuously, along with the complete target performance. When the score position given by the music tracker corresponds to a part in the score that, according to our estimated voice/music ratio, is dominated by voice(s), we initialize the accumulated cost vector of the lyrics tracker by values of $+\infty$ everywhere, and a value of 0 at the score position given by the music tracker. We then use separately music and lyrics trackers but we rely only on the score position given by the lyrics tracker. As soon as the score position given by the lyrics tracker corresponds to a part in the score dominated by music, we stop the lyrics tracker and rely on the position given by the music tracker.

4 Experiments and Discussion

For our experiments, we compare three different tracking models. The first, *music*, reproduces the work in (Brazier and Widmer, 2020b) and uses the music tracker only (including acoustic event detectors to deal with interludes and other unexpected events such as applause and acting pauses). The second one, *musicP*, is the state-of-the-art opera tracker (Brazier and Widmer, 2020a); it uses two music trackers in parallel, one using the features detailed in Section 3.1, the other using optimized audio features that have been tuned on the recitative subset of the Fischer performance. Finally, the third tracker *M&L* is the contribution of this paper. The systems are evaluated by their alignment accuracies (Cont et al., 2007). We report the mean error in ms, as well as the propor-

Conductor	Tracker	Mean	$\leq 1s$	$\leq 2s$	$\leq 5s$
Fischer	music	811ms	91.8%	95.0%	97.3%
	musicP	373ms	93.4%	96.8%	99.0%
	M&L	335ms	94.1%	97.3%	99.2%
Manacorda	music	561ms	90.1%	94.5%	97.9%
	musicP	547ms	90.3%	94.7%	98.0%
	M&L	410ms	91.6%	95.9%	99.0%

Table 2: Tracking error of three trackers: music (Brazier and Widmer, 2020b), musicP (Brazier and Widmer, 2020a), and Music and Lyrics (M&L).

tions of bar boundaries (which reflect the precision of our ground truth annotations) that are detected with an error less than 1, 2, and 5 seconds. The results are given in Table 2.

For both live target performances, the proposed *music & lyrics* tracker achieves the best accuracy, beating the *music* tracker, and also the *musicP* tracker whose features were tuned on the Fischer performance. The accuracy improvement on Fischer is relatively small, but no fine-tuning on features is done in our proposal. The improvements on Manacorda are more substantial, dropping the mean error to 410 ms and increasing all the 3 percentages by at least one point.

We tried to take into account the delay of the lyrics tracker, in adding an offset to the score position given by the tracker, but the best results were achieved in ignoring this delay.

5 Conclusion

We have presented a new state-of-the-art method for tracking full-length opera performances. The method makes use of an acoustic model that estimates the sung lyrics (phoneme probability vectors) over time. The final model combines lyrics and music information (without requiring the written lyrics as input) via two specific trackers. The combination helps to improve the tracking accuracy of the performance.

The proposed method requires a part segmentation of the reference performance. The beginnings and ends of each part are directly given by the manual bar annotations, useful to also handle structural mismatches in opera (Brazier and Widmer, 2021a). However, we plan to emancipate ourselves from the manual annotations with the development of a method that fully autonomously segments a piece.

Acknowledgments

The research is supported by the European Union under the EU's Horizon 2020 research and innovation programme, Marie Skłodowska-Curie grant agreement No. 765068 ("MIP-Frontiers"). The LIT AI Lab is supported by the Federal State of Upper Austria.

References

- Andreas Arzt, Harald Frostel, Thassilo Gadermaier, Martin Gasser, Maarten Grachten, and Gerhard Widmer. 2015. Artificial Intelligence in the Concertgebouw. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2424–2430, Buenos Aires, Argentina.
- Andreas Arzt and Gerhard Widmer. 2015. Real-Time Music Tracking Using Multiple Performances as a Reference. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 357–363, Málaga, Spain.
- Andreas Arzt, Gerhard Widmer, and Simon Dixon. 2008. Automatic Page Turning for Musicians via Real-Time Machine Listening. In *Proc. of the European Conference on Artificial Intelligence (ECAI)*, pages 241–245, Patras, Greece.
- Charles Brazier and Gerhard Widmer. 2020a. Addressing the Recitative Problem in Real-time Opera Tracking. In *Proc. of the Frontiers of Research in Speech and Music conference (FRSM)*, Silchar, India.
- Charles Brazier and Gerhard Widmer. 2020b. Towards Reliable Real-Time Opera Tracking: Combining Alignment with Audio Event Detectors to Increase Robustness. In *Proc. of the Sound and Music Computing Conference (SMC)*, pages 371–377, Turin, Italy.
- Charles Brazier and Gerhard Widmer. 2021a. Handling Structural Mismatches in Real-time Opera Tracking. In *Proc. of the European Signal Processing Conference (EUSIPCO)*, Dublin, Ireland.
- Charles Brazier and Gerhard Widmer. 2021b. On-Line Audio-to-Lyrics Alignment Based on a Reference Performance. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, Online.
- Arshia Cont, Diemo Schwarz, Norbert Schnell, and Christopher Raphael. 2007. Evaluation of Real-Time Audio-to-Score Alignment. In *International Symp. on Music Information Retrieval (ISMIR)*, pages 315–316, Vienna, Austria.
- Roger B Dannenberg. 1984. An On-Line Algorithm For Real-Time Accompaniment. In *Proc. of the International Computer Music Conference (ICMC)*, pages 193–198, Paris, France.
- Simon Dixon. 2005. An On-Line Time Warping Algorithm for Tracking Musical Performances. In *Proc. of International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1727–1728, Edinburgh, Scotland, UK.
- Thassilo Gadermaier and Gerhard Widmer. 2019. A Study of Annotation and Alignment Accuracy for Performance Comparison in Complex Orchestral Music. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 769–775, Delft, The Netherlands.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks. In *Proc. of the International Conference on Machine Learning (ICML)*, pages 369–376, Pittsburgh, Pennsylvania, USA.
- Khaled Koutini, Hamid Eghbal-zadeh, Matthias Dorfer, and Gerhard Widmer. 2019. The Receptive Field as a Regularizer in Deep Convolutional Neural Networks for Acoustic Scene Classification. In *Proc. of the European Signal Processing Conference (EUSIPCO)*, A Coruña, Spain.
- Olivier Lartillot, Carlos Cancino-Chacón, and Charles Brazier. 2020. Real-Time Visualisation of Fugue Played by a String Quartet. In *Proc. of the Sound and Music Computing Conference (SMC)*, pages 115–122, Turin, Italy.
- Gabriel Meseguer-Brocal, Alice Cohen-Hadria, and Geoffroy Peeters. 2018. DALI: A Large Dataset of Synchronized Audio, Lyrics and Notes, Automatically Created using Teacher-Student Machine Learning Paradigm. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 431–437, Paris, France.
- Matthew Prockup, David Grunberg, Alex Hrybyk, and Youngmoo E. Kim. 2013. Orchestral Performance Companion: Using Real-Time Audio to Score Alignment. *IEEE MultiMedia*, 20(2):52–60.
- Andrea Vaglio, Romain Hennequin, Manuel Mousallam, Gaël Richard, and Florence D'alché-Buc. 2020. Multilingual lyrics-to-audio alignment. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 512–519, Montréal, Canada.