FLEXIBLE SCORE FOLLOWING: THE PIANO MUSIC COMPANION AND BEYOND

Andreas Arzt^{1,2}, Werner Goebl³, Gerhard Widmer^{1,2}

¹Department of Computational Perception, Johannes Kepler University, Linz, Austria ²Austrian Research Institute for Artificial Intelligence (OFAI), Vienna, Austria ³Institute of Music Acoustics, University of Music and Performing Arts, Vienna, Austria andreas.arzt@jku.at

ABSTRACT

In our talk we will present a piano music companion that is able to follow and understand (at least to some extent) a live piano performance. Within a few seconds the system is able to identify the piece that is being played, and the position within the piece. It then tracks the performance over time via a robust score following algorithm. Furthermore, the system continuously re-evaluates its current position hypotheses within a database of scores and is capable of detecting arbitrary 'jumps' by the performer. The system can be of use in multiple ways, e.g. for piano rehearsal, for live visualisation of music, and for automatic page turning. At the conference, we will demonstrate this system live on stage. If possible, we would also like to encourage (hobby-)pianists in the audience to try the companion themselves. Additionally, we will give an outlook on our efforts to extend this approach to classical music in general, including heavily polyphonic orchestral music.

1. INTRODUCTION

Score following, also known as real-time music tracking, is a big challenge. It involves listening to a live incoming audio stream, extracting features from the audio that capture aspects of the 'sound' of the current moment, and tracking the most likely position in the score – regardless of the specific tempo chosen by the musicians on stage, of tempo changes due to expressive timing, and robust to varying sound quality and instrument sounds.

Real-time music tracking originated in the 1980s (see [1, 2]) and has attracted quite some research in recent years [3, 4, 5, 6, 7, 8, 9]. Although there still are many open research questions (such as on-line learning of predictive tempo models during a performance), score following is already being used in real-world applications. Examples include Antescofo¹, which is actively used by professional musicians to synchronise a performance (mostly solo instruments or small ensembles) with computer realised elements, and Tonara², a music tracking application focusing on the amateur pianist and running on the iPad.

In this paper we will summarise some recent developments of our music tracking system. First, in Section 2 below, we will describe a system that allows flexible music tracking on a big database of piano music. In Section 3 we will discuss some of the challenges we encountered when we prepared our music tracker to follow orchestral music live at a world-famous concert hall. Finally, in Section 4, we will discuss further steps of how to combine these two concepts – flexible score following on a database, and tracking orchestral music (or any other kind of instrumental classical music) –, and possible application enabled by this technology.

2. THE PIANO MUSIC COMPANION

The piano music companion is a versatile system that can be used by pianists, and more widely by consumers of piano music, in various scenarios. It is able to identify, follow and understand live performances of classical piano music – at least to some extent. The companion has two important capabilities that we believe such a system must possess: (1) automatically identifying the piece it is listening to, and (2) following the progress of the performer(s) within the score over time.

The companion is provided with a database of sheet music in symbolic form. Currently the database includes, amongst others, the complete solo piano works by Chopin and the complete Beethoven piano sonatas, and consists of roughly 1,000,000 notes in total. When listening to live music, the companion is able to identify the piece that is being played, and the position within the piece. It then tracks the progress of the performers over time, i.e. at any time the current position in the score is computed. Furthermore, it continuously re-evaluates its hypothesis and tries to match the current input stream to the complete database. Thus, it is able to follow any action of the musician, e.g. jumps to a different position or an entirely different piece – as long as the piece is part of the database. See Figure 1 for a very brief technical description. More information can be found in [10] and [11].

The piano music companion enables a range of applications, for both listeners and performers. Musicians might use this system as a convenient way to query and show the sheet music. They can just sit down at the piano, play a few notes, and the piano music companion will show the respective piece, at the correct position. The companion will take care of turning the pages at the appropriate times, and will also recognise jumps to a different positions or pieces and show the sheet music accordingly. For listeners of classical music, the main purpose of the piano music companion is to enrich the experience of classical music concerts. This might be achieved e.g. by providing information synchronised to the music, in various formats (text, images, videos). We will discuss this in a bit more detail in Section 3 below.

3. ADVENTURES IN THE CONCERTGEBOUW

So far, in our research we were focusing mainly on classical piano music, which resulted in the piano music companion briefly described above. The multi-national European research project PHENICX³ then provided us with the unique opportunity (and

¹repmus.ircam.fr/antescofo

 $^{^{2}}$ tonara.com

³http://phenicx.upf.edu



Figure 1: The piano music companion takes as input a live audio stream. It first tries to transcribe this stream into symbolic information (*Note Recogniser*). Then, it matches this information to a database of sheet music, also represented in symbolic form, via a fingerprinting algorithm (*Symbolic Music Matcher*). The output of this stage is fed to a *Multi Agent Music Tracking System* which is based on an on-line version of the dynamic time warping algorithm. This component tries to align the incoming audio stream to (synthesised versions of) the sheet music at the respective positions provided by the fingerprinter. At each point in time the best position (regarding alignment costs) is returned.

challenge) to work with the famous Concertgebouw Orchestra, and to demonstrate our score following technology in the context of a big, real-life symphonic concert.

The event took place on February 7th, 2015, in the Concertgebouw in Amsterdam. The Royal Concertgebouw Orchestra, conducted by Semyon Bychkov, performed the *Alpensinfonie* (Alpine Symphony) by Richard Strauss. This concert was part of a series called 'Essentials', during which technology developed within the project can be tested in a real-life concert environment. For the demonstration a test audience of about 30 people was provided with tablet computers and placed in the rear part of the concert hall. The music tracker was used to control the transmission and display of additional information, synchronised to the live performance on stage (a similar study was presented in [12]). The user interface and the visualisations were provided by our project partner Videodock⁴.



Figure 2: The Multi-agent Tracker. The live input is fed to N independent instances of the tracking algorithm. Each aligns the input to its own score representation, based on different performances of the same piece. Then, the individual hypotheses are combined and the estimate of the current position in the score is returned.

During the concert the audience was presented with 3 different kinds of visualisations. The sheet music was shown synchronised to the live performance, with highlighting of the cur-

⁴http://videodock.com

rent bar and automatic turning of the pages. Textual information, prepared by a musicologist was shown at the appropriate times, helping the audience to understand the structure of the piece. Additionally, artistic videos were shown, which were telling the story of the *Alpensinfonie*.

As this was a live experiment during a real concert, our main goal during the preparation was to make sure that the algorithm will not get lost at any point in time. In the end, this led to a method that increased both the robustness and also the accuracy of the tracking process, with the main idea being to use multiple actual performances (which are preprocessed automatically) as the basis of the tracking process, instead of a single synthesised version of the score (see Figure 2 for a brief explanation). For further information about the experiment at the Concertgebouw we refer the reader to [13], and to [14] for a larger scale evaluation of the multi -gent tracking system.

4. OUTLOOK: THE COMPLETE CLASSICAL MUSIC COMPANION

In the previous two sections above we introduced a flexible piano music tracker and a robust tracking algorithm for orchestral music – and basically any other kind of classical music. Our vision is to finally combine the capabilities of these two systems, resulting in a prototype of what we call the *Complete Classical Music Companion*.

In the end, we envisage a mobile application on a tablet computer or a mobile phone that is at your fingertips anytime and anywhere and can act as your personal classical music companion. Always ready to listen to classical music, and trying to identify and to understand what it is listening to, it can provide you with all kinds of information, guide you, and help you understand the performance. It can provide you with background information about the composer and the historical background of the piece. During the performance it can present information synchronised to the live performance, provided by musicologists and adapted to your level of expertise. It can also show you the sheet music or more abstract representations of the piece (e.g. regarding the structure or the instrumentation).

In the future, we are planning not to limit ourselves to information about the piece, but would also like to include live analysis of the performance itself. As a first step, we will try to build an on-line version of the visualisation presented in [15], which tries to capture two important aspects of an expressive performance – tempo and loudness – and visualises them in a 2-dimensional space.

The biggest obstacle right now is that one important component, the *Note Recogniser* (see Figure 1), is limited to piano music only. Already the transcription of polyphonic piano music to a symbolic representation (on-line!) is a very difficult task, and so far we do not have a solution that works sufficiently well for e.g. orchestral music.

Another big challenge still is the data preparation for the companion. For every single piece both a symbolic version (e.g. MusicXML) and images of the sheet music – synchronised to each other – are needed. Unfortunately, optical music recognition programs are far from being capable of automatically transforming sheet music into MusicXML in sufficient quality, especially for complicated orchestral scores. Currently we are slowly adding new pieces, which involves a lot of manual effort, but in the long run a different solution is needed – either crowd based or in cooperation with a music publisher.

5. CONCLUSIONS

In this paper we summarised our efforts towards building a *Complete Classical Music Companion*. We already have a working prototype for piano music only, which we will demonstrate during the conference. Also, we have successfully demonstrated that our music tracker is well-suited to follow orchestral performances under real-life conditions.

There are still a few parts missing to enable us to finally build a prototype of the system we have in mind. Main future work will be concentrated on filling these gaps, the most important one being to find a way to make the transcription component work with a wider range of instruments.

6. ACKNOWLEDGEMENTS

This research is supported by the Austrian Science Fund (FWF) under project number Z159 and the EU FP7 Project PHENICX (grant no. 601166).

7. REFERENCES

- [1] Roger Dannenberg, "An on-line algorithm for real-time accompaniment," in *Proc. of the International Computer Music Conference (ICMC)*, Paris, France, 1984.
- [2] Barry Vercoe, "The synthetic performer in the context of live performance," in *Proc. of the International Computer Music Conference (ICMC)*, Paris, France, 1984.
- [3] Andreas Arzt, Gerhard Widmer, and Simon Dixon, "Automatic page turning for musicians via real-time machine listening," in *Proc. of the European Conference on Artificial Intelligence (ECAI)*, Patras, Greece, 2008.
- [4] Arshia Cont, "A coupled duration-focused architecture for realtime music to score alignment," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 6, pp. 837–846, 2009.
- [5] Christopher Raphael, "Music Plus One and machine learning," in Proc. of the International Conference on Machine Learning (ICML), Haifa, Israel, 2010.
- [6] Takuma Otsuka, Kazuhiro Nakadai, Toru Takahashi, Tetsuya Ogata, and Hiroshi G. Okuno, "Real-time audio-toscore alignment using particle filter for co-player music robots," *EURASIP Journal on Advances in Signal Processing*, vol. 2011, no. 2011:384651, 2011.
- [7] Nicola Montecchio and Arshia Cont, "A unified approach to real time audio-to-score and audio-to-audio alignment using sequential montecarlo inference techniques," in *Proc. of the IEEE Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Prague, Czech Republic, 2011.
- [8] Zhiyao Duan and Bryan Pardo, "A state space model for on-line polyphonic audio-score alignment," in *Proc. of the IEEE Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Prague, Czech Republic, 2011.
- [9] Filip Korzeniowski, Florian Krebs, Andreas Arzt, and Gerhard Widmer, "Tracking rests and tempo changes: Improved score following with particle filters," in *Proc.* of the International Computer Music Conference (ICMC), Perth, Australia, 2013.

- [10] Andreas Arzt, Sebastian Böck, and Gerhard Widmer, "Fast identification of piece and score position via symbolic fingerprinting," in *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, Taipeh, Taiwan, 2012.
- [11] Andreas Arzt, Sebastian Böck, Sebastian Flossmann, Harald Frostel, Martin Gasser, and Gerhard Widmer, "The complete classical music companion v0.9," in *Proc. of the AES Conference on Semantic Audio*, London, England, 2014.
- [12] Matthew Prockup, David Grunberg, Alex Hrybyk, and Youngmoo E. Kim, "Orchestral performance companion: Using real-time audio to score alignment," *IEEE Multimedia*, vol. 20, no. 2, pp. 52–60, 2013.
- [13] Andreas Arzt, Harald Frostel, Thassilo Gadermaier, Martin Gasser, Maarten Grachten, and Gerhard Widmer, "Artificial intelligence in the concertgebouw," in *Proc. of the International Joint Conference on Artificial Intelligence* (*IJCAI*), Buenos Aires, Argentina, 2015.
- [14] Gerhard Arzt, Andreas Widmer, "Real-time music tracking using multiple performances as a reference," in *Proc.* of the International Society for Music Information Retrieval Conference (ISMIR), Malaga, Spain, 2015.
- [15] Jörg Langner and Werner Goebl, "Visualizing expressive performance in tempo-loudness space," *Computer Music Journal*, vol. 27, 2003.