The Piano Music Companion

Andreas Arzt⁽¹⁾, Sebastian Böck⁽¹⁾, Sebastian Flossmann⁽¹⁾, Harald Frostel⁽¹⁾, Martin Gasser⁽²⁾, Cynthia C. S. Liem⁽³⁾ and Gerhard Widmer^{(1,2)1}

Abstract. We present a system that we call 'The Piano Music Companion' and that is able to follow and understand (at least to some extent) a live piano performance. Within a few seconds this system can identify the piece that is being played, and the position within the piece. It then tracks the progress of the performer over time via a robust score following algorithm. The companion is useful in multiple ways, e.g., it can be used for piece identification, music visualisation, during piano rehearsal and for automatic page turning.

1 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.

To support these two capabilities, the companion is provided with a database of sheet music in symbolic form, i.e., sequences of (noteon, pitch) pairs. 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 (about 330 pieces). 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. The system is tolerant to performance errors and slight variations, and is robust to tempo changes. These capabilities enable various application, of which four are described in Sec. 2 below.

2 Applications

The companion can be used to *identify classical piano music*. Wherever you are, and whatever the source of music, be it a live concert, a DVD or radio, only a few seconds of audio material are required to confidently identify the piece and retrieve the name, the composer, and additional meta-data like the historical context of the piece, famous interpretations, and where to buy recordings. We want to emphasise that this task differs from audio identification as provided for popular music by services like Shazam [5]. Given a query, these services are able to identify exact copies of the audio in their database, i.e., instances of exactly the same performance of the piece. In contrast to this, we are interested in the piece, i.e., the composition that was the basis for the performance. Hence, the companion should not rely on having *every* available performance as an audio file in the database, but on a more general representation: the sheet music. Furthermore, this also ensures that the companion also works for (so far unknown) live performances.

The second use case is *live music visualisation and performance enrichment*. As, at any point in time, the companion knows exactly where in the sheet music the performer is, it can show visualisations synchronised to the live music. In the simplest case it can show the sheet music itself, with a marker showing the current position. While this is already helpful for listeners, more sophisticated visualisations and enrichments are possible, like showing information about the structure of the piece and the most important themes, and giving hints about what to listen for at specific moments.

A third application is the use of the companion during *piano re-hearsal*, as the system can follow a performer and show the sheet music accordingly, even if the he/she repeats a section over and over or only plays parts of the score he/she needs to rehearse. The musician can simply sit down at the piano, start the app on a tablet computer, query the piece by playing the first few beats and start practising. The companion will follow all the actions and show the sheet music on the screen (which has the additional benefit of not having to carry heavy books to practise sessions).



Figure 1. The fully automatic page turner in action

Finally, professional musicians can use the companion, on-stage, for *fully automatic page turning*. Being able to track the live performance and thus at all times knowing where the musicians currently are in the piece, the system can control and automatically trigger a mechanical page turning device that turns the sheet music page at the appropriate time. This way, musicians do not have to rely on a human page turner (who will always have to get in between the musician and

¹ (1) Johannes Kepler University, Linz, Austria; ⁽²⁾ Austrian Research Institute for Artificial Intelligence, Vienna, Austria; ⁽³⁾ Delft University of Technology, Delft, The Netherlands

the sheet music), nor risk taking their hands off the instrument to turn the pages themselves. At the same time, an automatic page turner allows them to use the paper version of the score, which they normally prefer. The automatic page turner has already been used in real piano recitals in the context of various scientific gala events in Vienna. The device in action is shown in Fig. 1.

3 System Overview

Fig. 2 gives an overview of the piano companion. As described above, the system is able to automatically detect the played piece and then track it over time. Two main components running in parallel enable the companion to achieve this: (1) an 'Instant Piece and Position Recogniser' and (2) a 'Multi Agent Music Tracking System'.

The recogniser is based on an on-line piano music transcription algorithm [3], which takes the audio stream and translates it into symbolic information (a list of pitches with time stamps) using a bidirectional recurrent neural network. The most recently detected notes of the live performance are then matched to the database of sheet music via a tempo-independent fingerprinting method [1]. This process is continuously running in the background, regularly providing new piece and (rough) position hypotheses for the tracking component.



Figure 2. System Overview

These hypotheses are then processed by a multi agent music tracking algorithm (for more information about the tracking algorithm see [2]), which tries to match the current audio input to these respective positions in the sheet music and follows the progress of the musician over time. At each point in time, a single tracker is marked as active, i.e., represents our system's belief about the current position in the database.

4 Conclusions and Future Work

In this paper we presented a piano music companion that is based on a very flexible music tracking algorithm. We see this system as a first step towards our vision of a "Complete Classical Music Companion". Here, we are thinking about a system that is at your fingertips anytime and anywhere, possibly as an app on a mobile device like a tablet computer, and that provides you with information about what is going on musically around you. Whatever piece, for whatever instrumentation, and whoever the performers are, the companion will inform you about both the written music and the specifics of the ongoing (live) performance, and guide you in the listening process.

An important step towards this goal is to lift the restriction to piano music only. We already have some encouraging preliminary results regarding tracking, i.e., we can track live performances by symphonic orchestras, at least well enough for certain applications like synchronised visualisation of the sheet music. Thus, the limiting component currently is the transcription system. Future work has to be done to find a solution that is both fast and accurate enough for our intended application.

All this is further motivated by the fact that the companion will play a practical role in the PHENICX² project [4], which has the broad goal of "changing the way we experience classical music concerts", with one of the project partners being the world-renowned Royal Concertgebouw Orchestra Amsterdam. Thus, future work will also include finding ways of using the companion in the practical context of a world-class orchestra.

ACKNOWLEDGEMENTS

This research is supported by the Austrian Science Fund (FWF) under project numbers Z159, TRP 109-N23 and the EU FP7 Project PHENICX (grant no. 601166).

REFERENCES

- Andreas Arzt, Gerhard Widmer, Sebastian Böck, Reinhard Sonnleitner, and Harald Frostel, 'Towards a complete classical music companion', in *Proceedings of the European Conference on Artificial Intelligence* (ECAI), (2012).
- [2] Andreas Arzt, Gerhard Widmer, and Simon Dixon, 'Automatic page turning for musicians via real-time machine listening', in *Proceedings of* the 18th European Conference on Artificial Intelligence (ECAI), (2008).
- [3] Sebastian Böck and Markus Schedl, 'Polyphonic piano note transcription with recurrent neural networks', in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, (2012).
- [4] Emilia Gómez, Maarten Grachten, Alan Hanjalic, Jordi Janer, Sergi Jorda, Carles F. Julia, Cynthia Liem, Agustin Martorell, Markus Schedl, and Gerhard Widmer, 'PHENICX: Performances as Highly Enriched aNd Interactive Concert experiences', in *Proceedings of the Sound and Music Computing Conference (SMC)*, (2013).
- [5] Avery Wang, 'An industrial strength audio search algorithm', in Proceedings of the International Conference on Music Information Retrieval (ISMIR), (2003).

² http://phenicx.upf.edu