SCORE-PERFORMANCE MATCHING

JOHN ASHLEY BURGOYNE · ID 260172518

MUMT 611 · PROF. ICHIRO FUJINAGA · 1 MARCH 2007

SCORE-PERFORMANCE MATCHING, the problem of teaching machines to follow along with live musical performers, has enjoyed nearly 25 years of study and development. It benefits composers of electro-acoustic music and classical musicians who need a rehearsal partner. With the advent of probabilistic models, score-performance matching can now be accomplished with good accuracy, and future work will look to improve sensitivity to the many parameters of expressive musical performances.

The primary challenges in score-performance matching are common to most 'live' systems: reducing latency to an imperceptible amount and running all computation in real time. The motivations for overcoming these challenges are manifold. Automatic accompaniment systems, which presuppose a reliable means of score-performance matching, can help soloists rehearse with virtual orchestras for much longer than is logistically possible with human ones. Performers in electro-acoustic contexts, often confused by foot pedals and other triggering devices, can focus on their instruments instead. Digital music stands could render the infamous arguments over who should be the page-turner obsolete and spare percussionists the error-prone task of counting extended rests.

HISTORY OF THE PROBLEM

The field of score-performance matching opened in Paris, France, at the International Computer Music Conference of 1984. Two papers, one by Barry Vercoe and another by Roger Dannenberg, presented alternative

systems for the task (Vercoe 1984; Dannenberg 1984). Vercoe's paper was quite short but succeeded in outlining the key sub-components of any algorithm for following scores with the ultimate goal of accompanying a live performer: listening, performing, and learning. Dannenberg's paper, considerably longer, detailed a string-matching algorithm for the task. Although in keeping with the trends in artificial intelligence and algorithmics at its time, this paper did not include an awareness of tempo or probabilistic models, which have become important features of successful systems recently.

Since Vercoe and Dannenberg, work in score-performance matching has been continual and attracted well-known researchers in computer music (Puckette 1995; Desain, Honing, and Heijink 1997; Orio, Lemouton, and Schwarz 2003). Dannenberg himself has continued working on the problem, most recently in collaboration with Lorin Grubb. The Grubb-Dannenberg system estimates a continuous probability distribution function over possible future onsets to track the location of vocal performers in real time (Grubb and Dannenberg 1997; Grubb and Dannenberg 1998). Other successful systems in recent research, including the system currently in use at IRCAM, have been based on hidden Markov models (HMMs), a more advanced probabilistic model (Cano, Loscos, and Bonada 1999; Schwarz, Cont, and Schnell 2005). Probabilistic systems like these offer a good compromise between having an awareness of tempo and offering the soloist flexibility to change it, and due to their inferential properties, they also help to reduce latency. Whatever the model, most new advances in the field are presented at the International Computer Music Conference.

MUSIC PLUS ONE

Music Plus One is another score-performance matching system based on HMMs, but its development of the basic model is more elaborate than its competitors. Like the Vercoe model, it is divided into phases: 'Listen'

and 'Play'. Given a complete score of the solo part, 'Listen' extracts onset estimates from the soloist in real time (Raphael 1999). Given these onset estimates and a full score, 'Play' realises an accompaniment in real time while the soloist is performing (Raphael 2001). Inspired by old Music Minus One recordings, the system focuses on synthesising live accompaniments for common-practise solo concertos, although its techniques could be extended to other genres. Although the mathematical details behind the system are heavy, most have simple musical interpretations, e.g., a degree of *rubato* and how it might shift over time relative to a 'carrier' tempo than might be changing more gradually. Originally oriented toward MIDI-based accompaniments, Music Plus One has also been extended to use time-scaled audio recordings of concerto accompaniments (Raphael 2003).

Music Plus One is astonishingly good at its intended task, but open questions remain in the field of score-performance matching. Most importantly, all current systems consider only the timing aspects of accompaniment, i.e., how to ensure that the accompaniment plays at its intended time. It does not, however, respond to changes of dynamic or other expressive gestures that are commonly believed to be an essential aspect of collaborative performances among human performers. Moreover, all current systems assume that a single soloist must be tracked, not a polyphonic ensemble. An automatic electronic addition to an orchestra, for example, would require some some mechanism of polyphonic pitch tracking keep its place, which is not yet a solved problem. The next 25 years promise more exciting work.

REFERENCES

Cano, P., A. Loscos, and J. Bonada. 1999. Score-performance matching using HMMs. In *Proceedings of the International Computer Music Conference*.

Dannenberg, R. B. 1984. An on-line algorithm for real-time accompa-

niment. In *Proceedings of the International Computer Music Conference*, Paris, France, 193–98.

Desain, P., H. Honing, and H. Heijink. 1997. Robust score-performance matching: Taking advantage of structural information. In *Proceedings of the International Computer Music Conference*, 337–40.

Grubb, L., and R. B. Dannenberg. 1997. A stochastic method for tracking a vocal performer. In *Proceedings of the International Computer Music Conference*, 301–08.

Grubb, L., and R. B. Dannenberg. 1998. Enhanced vocal performance tracking using multiple information sources. In *Proceedings of the International Computer Music Conference*, 37–44.

Orio, N., S. Lemouton, and D. Schwarz. 2003. Score following: State of the art and new developments. In *Proceedings of the International Conference on New Interfaces for Musical Expression*.

Puckette, M. 1995. Score following using the sung voice. In *Proceedings of the International Computer Music Conference*.

Raphael, C. 1999. Automatic segmentation of acoustic musical signals using hidden Markov models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21 (4): 360–70.

Raphael, C. 2001. A probabilistic system for automatic musical accompaniment. *Journal of Computational and Graphical Statistics* 10 (3): 487–512.

Raphael, C. 2003. Orchestral musical accompaniment from synthesized audio. In *Proceedings of the International Computer Music Conference*.

Schwarz, D., A. Cont, and N. Schnell. 2005. From Boulez to ballads: Training IRCAM's score follower. In *Proceedings of the International Computer Music Conference*, Barcelona, Spain.

Vercoe, B. 1984. The synthetic performer in the context of live performance. In *Proceedings of the International Computer Music Conference*, Paris, France, 199–200.