

# Expressivity-Preserving Tempo Transformation for Music – A Case-Based Approach

Ramon López de Mántaras, Maarten Grachten, and Josep-Lluís Arcos

IIIA, Artificial Intelligence Research Institute  
CSIC, Spanish Council for Scientific Research  
Campus UAB, 08193 Bellaterra, Catalonia, Spain  
Vox: +34-93-5809570; Fax: +34-93-5809661  
{mantaras,arcos,maarten}@iia.csic.es

**Abstract.** The research described in this paper focuses on global tempo transformations of monophonic audio recordings of saxophone jazz performances. More concretely, we have investigated the problem of how a performance played at a particular tempo can be automatically rendered at another tempo while preserving its expressivity. To do so we have developed a case-based reasoning system called *TempoExpress*. The results we have obtained have been extensively compared against a standard technique called uniform time stretching (UTS), and show that our approach is superior to UTS.

## 1 The Problem of Generating Expressive Music

It has been long established that when humans perform music from score, the result is never a literal, mechanical rendering of the score (the so-called nominal performance). As far as performance deviations are intentional (that is, they originate from cognitive and affective sources as opposed to e.g. motor sources), they are commonly thought of as conveying musical expressivity, which forms an important aspect of music. Two main functions of musical expressivity are generally recognized. Firstly, expressivity is used to clarify the musical structure (in the broad sense of the word: this includes for example metrical structure [Sloboda, 1983], but also the phrasing of a musical piece [Gabrielsson, 1987], and harmonic structure [Palmer, 1996]). Secondly, expressivity is used as a way of communicating, or accentuating affective content [Juslin, 2001; Gabrielsson, 1995].

An important issue when performing music is the effect of tempo on expressivity. It has been argued that temporal aspects of performance scale uniformly when tempo changes [Repp, 1994]. That is, the durations of all performed notes maintain their relative proportions. This hypothesis is called relational invariance (of timing under tempo changes). However, counter-evidence for this hypothesis has been provided [Desain and Honing, 1994; Timmers et al., 2002], and a recent study shows that listeners are able to determine above chance-level whether audio-recordings of jazz and classical performances are uniformly time stretched or original recordings, based solely on expressive aspects of the performances [Honing, 2006]. Our approach also experimentally refutes the relational

invariance hypothesis by comparing the automatic transformations generated by *TempoExpress* against uniform time stretching.

## 2 TempoExpress

Given a MIDI score of a phrase from a jazz standard, and given a monophonic audio recording of a saxophone performance of that phrase at a particular tempo (the source tempo), and given a number specifying the target tempo, the task of the system is to render the audio recording at the target tempo, adjusting the expressive parameters of the performance to be in accordance with that tempo.

*TempoExpress* solves tempo transformation problems by case-based reasoning. Problem solving in case-based reasoning is achieved by identifying and retrieving a problem (or a set of problems) most similar to the problem that is to be solved from a case base of previously solved problems (also called cases), and adapting the corresponding solution to construct the solution for the current problem.

To realize a tempo transformation of an audio recording of an input performance, *TempoExpress* needs an XML file containing the melodic description of the recorded audio performance, a MIDI file specifying the score, and the target tempo to which the performance should be transformed (the tempo is specified in the number of beats per minute, or BPM). The result of the tempo transformation is an XML file containing the modified melodic description, that is used as the basis for synthesis of the transformed performance. For the audio analysis (that generates the XML file containing the melodic description of the input audio performance) and for the audio synthesis, *TempoExpress* relies on an external system for melodic content extraction from audio, developed by Gómez et al. [2003b]. This system performs pitch and onset detection to generate a melodic description of the recorded audio performance, the format of which complies with an extension of the MPEG7 standard for multimedia content description [Gómez et al., 2003a].

We apply the edit-distance [Levenshtein, 1966] in the retrieval step in order to assess the similarity between the cases in the case base (human performed jazz phrases at different tempos) and the input performance whose tempo has to be transformed. To do so, firstly the cases whose performances are all at tempos very different from the source tempo are filtered out. Secondly, the cases with phrases that are melodically similar to the input performance (according to the edit-distance) are retrieved from the case base. The melodic similarity measure we have developed for this is based on abstract representations of the melody [Grachten et al., 2005] and has recently won a contest for symbolic melodic similarity computation (MIREX 2005).

In the reuse step, a solution is generated based on the retrieved cases. In order to increase the utility of the retrieved material, the retrieved phrases are split into smaller segments using a melodic segmentation algorithm [Temperley, 2001]. As a result, it is not necessary for the input phrase and the retrieved phrase to match as a whole. Instead, matching segments can be reused from various retrieved phrases. This leads to the generation of *partial* solutions for the

input problem. To obtain the complete solution, we apply *constructive adaptation* [Plaza and Arcos, 2002], a reuse technique that constructs complete solutions by searching the space of partial solutions.

The solution of a tempo-transformation consists in a performance annotation. This performance annotation is a sequence of changes that must be applied to the score in order to render the score expressively. The result of applying these transformations is a sequence of performed notes, the output performance, which can be directly translated to a melodic description at the target tempo, suitable to be used as a directive to synthesize audio for the transformed performance.

To our knowledge, all current performance rendering systems deal with predicting expressive values like timing and dynamics for the notes in the score. Contrastingly, *TempoExpress* not only predicts values for timing and dynamics, but also deals with more extensive forms of musical expressivity, like note insertions, deletions, consolidations, fragmentations, and ornamentations.

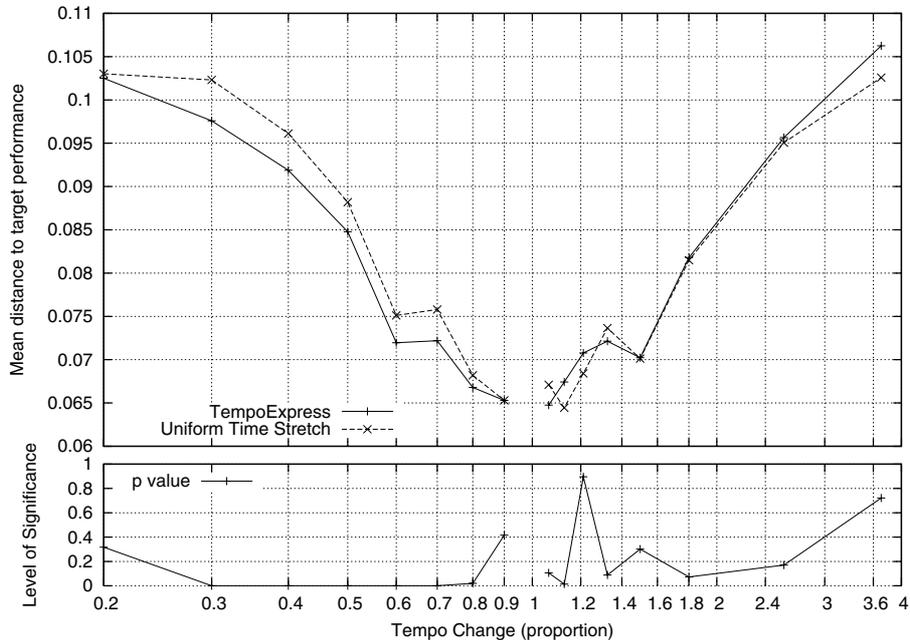
### 3 Results

In this section we describe results of an extensive comparison of *TempoExpress* against *uniform time stretching* (UTS), the standard technique for changing the tempo of audio recordings, in which the temporal aspects (such as note durations and timings) of the recording are scaled by a constant factor proportional to the tempo change.

For a given tempo transformation task, the correct solution is available as a target performance: a performance at the target tempo by a professional musician, that is known to have appropriate expressive values for that tempo. The results of both tempo transformation approaches are evaluated by comparing them to the target performance. More specifically, let  $M_H^s$  be a melodic description of a performance of phrase  $p$  by a musician  $H$  at the source tempo  $s$ , and let  $M_H^t$  be a melodic description of a performance of  $p$  at the target tempo  $t$  by  $H$ . Using *TempoExpress* (TE), and UTS, we derive two melodic descriptions for the target tempo from  $M_H^s$ , respectively  $M_{TE}^t$ , and  $M_{UTS}^t$ .

We evaluate both derived descriptions by their similarity to the target description  $M_H^t$ . To compute the similarity we use a distance measure that has been modeled after human perceived similarity between musical performances. Ground truth for this was gathered through a web-survey in which human subjects rated the perceived dissimilarity between different performances of the same melodic fragment. The results of the survey were used to optimize the parameters of an edit-distance function for comparing melodic descriptions. The optimized distance function correctly predicts 85% of the survey responses.

In this way, the results of *TempoExpress* and UTS were compared for 6364 tempo-transformation problems, using 64 different melodic segments from 14 different phrases. The results are shown in figure 1. The figure shows the distance of both *TempoExpress* and UTS results to the target performances, as a function of tempo change (measured as the ratio of the target tempo to the source tempo). The lower plots show the significance value for the null hypothesis



**Fig. 1.** Performance of *TempoExpress* vs. UTS as a function of the ratio of target tempo to source tempo. The lower plot shows the probability of incorrectly rejecting  $H_0$  for the Wilcoxon signed-rank tests.

that the melodic descriptions generated by *TempoExpress* are not more similar or less similar to the target description than the melodic description generated using UTS (in other words, the hypothesis that *TempoExpress* does not give an improvement over UTS).

Firstly, observe that the plot in Figure 1 shows an increasing distance to the target performance with increasing tempo change (both for slowing down and for speeding up), for both tempo transformation techniques. This is evidence against the hypothesis of relational invariance discussed earlier in this paper. This hypothesis implies that the UTS curve should be horizontal, since under relational variance, tempo transformations are supposed to be achieved through mere uniform time stretching.

Secondly, a remarkable effect can be observed in the behavior of *TempoExpress* with respect to UTS, which is that *TempoExpress* improves the result of tempo transformation specially when slowing performances down. When speeding up, the distance to the target performance stays around the same level as with UTS. In the case of slowing down, the improvement with respect to UTS is mostly significant, as can be observed from the lower part of the plot. Note that the p-values are rather high for tempo change ratios close to 1, meaning that for those tempo changes, the difference between *TempoExpress* and UTS is not statistically significant. This is in accordance with the common sense that

**Table 1.** Overall comparison between *TempoExpress* and uniform time stretching, for upwards and downwards tempo transformations, respectively

	mean distance to target		Wilcoxon signed-rank test		
	<i>TempoExpress</i>	UTS	p <>	z	df
tempo ↑	0.0791	0.0785	0.046	1.992	3181
tempo ↓	0.0760	0.0786	0.000	9.628	3181

slight tempo changes do not require many changes, in other words, relational invariance approximately holds when the amount of tempo change is very small.

Table 1 summarizes the results for both tempo increase and decrease. Columns 2 and 3 show the average distance to the target performance for *TempoExpress* and UTS, averaged over all tempo increase problems, and tempo decrease problems respectively. The other columns show data from the Wilcoxon signed-rank test. The p-values are the probability of incorrectly rejecting  $H_0$  (that there is no difference between the *TempoExpress* and UTS results). This table also shows that for downward tempo transformations, the improvement of *TempoExpress* over UTS is small, but extremely significant ( $p < .001$ ), whereas for upward tempo transformations UTS seems to be better, but the results are slightly less decisive ( $p < .05$ ).

## 4 Conclusions

In this paper we have summarized our research results on a case-based reasoning approach to global tempo transformations of music performances, focusing on saxophone recordings of jazz themes. We have addressed the problem of how a performance played at a particular tempo can be automatically rendered at another tempo preserving expressivity. Moreover, we have described the results of an extensive experimentation over a case-base of more than six thousand transformation problems. *TempoExpress* clearly performs better than UTS when the target problem is slower than the source tempo. When the target tempo is higher than the source tempo the improvement is less significant. Nevertheless, *TempoExpress* behaves as UTS except in transformations to very fast tempos. This result may be explained by a lack of example cases with fast tempos.

## Bibliography

- Desain, P. and Honing, H. (1994). Does expressive timing in music performance scale proportionally with tempo? *Psychological Research*, 56:285–292.
- Gabrielsson, A. (1987). Once again: The theme from Mozart’s piano sonata in A major (K. 331). A comparison of five performances. In Gabrielsson, A., editor, *Action and perception in rhythm and music*, pages 81–103. Royal Swedish Academy of Music, Stockholm.
- Gabrielsson, A. (1995). Expressive intention and performance. In Steinberg, R., editor, *Music and the Mind Machine*, pages 35–47. Springer-Verlag, Berlin.

- Gómez, E., Gouyon, F., Herrera, P., and Amatriain, X. (2003a). Using and enhancing the current MPEG-7 standard for a music content processing tool. In *Proceedings of Audio Engineering Society, 114th Convention*, Amsterdam, The Netherlands.
- Gómez, E., Grachten, M., Amatriain, X., and Arcos, J. L. (2003b). Melodic characterization of monophonic recordings for expressive tempo transformations. In *Proceedings of Stockholm Music Acoustics Conference 2003*.
- Grachten, M., Arcos, J. L., and López de Mántaras, R. (2005). Melody retrieval using the Implication/Realization model. MIREX <http://www.music-ir.org/evaluation/mirex-results/articles/similarity/grachten.pdf>.
- Honing, H. (2006). Is expressive timing relational invariant under tempo transformation? *Psychology of Music*. (in press).
- Juslin, P. (2001). Communicating emotion in music performance: a review and a theoretical framework. In Juslin, P. and Sloboda, J., editors, *Music and emotion: theory and research*, pages 309–337. Oxford University Press, New York.
- Levenshtein, V. I. (1966). Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady*, 10:707–710.
- Palmer, C. (1996). Anatomy of a performance: Sources of musical expression. *Music Perception*, 13(3):433–453.
- Plaza, E. and Arcos, J. L. (2002). Constructive adaptation. In Craw, S. and Preece, A., editors, *Advances in Case-Based Reasoning*, number 2416 in Lecture Notes in Artificial Intelligence, pages 306–320. Springer-Verlag.
- Repp, B. H. (1994). Relational invariance of expressive microstructure across global tempo changes in music performance: An exploratory study. *Psychological Research*, 56:285–292.
- Sloboda, J. A. (1983). The communication of musical metre in piano performance. *Quarterly Journal of Experimental Psychology*, 35A:377–396.
- Temperley, D. (2001). *The Cognition of Basic Musical Structures*. MIT Press, Cambridge, Mass.
- Timmers, R. and Ashley, R., Desain, P., Honing, H., and Windsor, L. (2002). Timing of ornaments in the theme of Beethoven’s Paisiello Variations: Empirical data and a model. *Music Perception*, 20(1):3–33.