MIRTEMPO 1.8: TEMPO ESTIMATION BY TRACKING A COMPLETE METRICAL STRUCTURE

Olivier Lartillot
University of Oslo
olivier.lartillot@imv.uio.no

ABSTRACT

This paper describes the two variants OL1 and OL2 of the model submitted to the MIREX 2018 tempo estimation tasks, and compare them with respect to my previous submission for MIREX 2013 [3].

1. INTRODUCTION

The tempo estimation model has been built within the *MIRtoolbox* platform [1] [1]. We submitted a previous version of the method to MIREX 2013 [3], and integrated it in version 1.6 of *MIRtoolbox* [2]. The new improvements presented in this paper will be made available in the upcoming version 1.8 of *MIRtoolbox*.

2. ONSET CURVE

We use our *'Emerge'* onset detector released in *MIRtoolbox* 1.5 that can handle vibrato and dense textures [2, 3]. The method is based on an improvement and generalization of the flux method that look at particular time / frequency region and can tolerate spectral fluctuations of limited frequency range.

Instead of using a frequency resolution of at least 0.1 Hz, as in our previous submission [3], we can simply use a resolution of at least 1 Hz. This makes the computation significantly faster and less greedy in memory.

3. PERIODICITY ESTIMATION

Periodicity estimation is carried out using exactly the same principles as in our previous submission [3].

Tempo is estimated by computing autocorrelation functions, on a moving window of frame length 5 seconds and hop factor 5%, for a range of time lags that corresponds to a tempo range between 24 and 500 BPM. The autocorrelation curve is normalized so that the autocorrelation at zero lag is identically 1.

[1] http://www.jyu.fi/hum/laitokset/musiikki/en/research/coe/materials /mirtoolbox

[2] Due to a bug, this method was not working correctly in version 1.7.

One interesting problem with autocorrelation functions is that a lag can be selected as prominent because it is found often in the signal although the lag is not repeated successively. We propose a simple solution based on the following property: For a given lag to be repeated at least twice, the periodicity score associated with twice the lag should have a high probability score as well. This heuristics can be implemented as a single post-processing operations applied to the autocorrelation function, removing all periodicity candidate that to not have stronger periodicity at twice its lag.

4. PEAK PICKING

Peak picking is carried out using the same principles as in our previous submission [3].

It is applied to the frame-by-frame autocorrelation functions. The beginning and the end of the autocorrelation curves are not taken into consideration for peak picking as they do not correspond to actual local maxima. The only modification in the new submission is that the first peak—i.e., the one with lowest autocorrelation lag—should be preceded by a valley with negative autocorrelation. This enables to filter out non-relevant peaks.

A given local maximum will be considered as a peak if its distance with the previous and successive local minima (if any) is higher than this threshold .05. This distance is expressed with respect to the total amplitude of the input signal. This distance of .05 is hence equivalent to 5 % of the distance between the global maximum and the minimum of the input signal [4]. The peak position and amplitude are estimated more precisely using quadratic interpolation.

5. TRACKING THE WHOLE METRICAL HIERARCHY

In the presence of a given pulsation in the musical excerpt that is being analyzed – let's say with a BPM of 120, i.e., with two pulses per second – the periodicity function will indicate a high periodicity score related to the period .5 s. But generally if there is a pulsation at a given tempo, multiples of the pulsation can also be found that are twice slower (1 s), three times slower, etc. For that reason, the periodicity function usually shows a series of peaks equally distant for all multiples of a given period. This has close connections with the notion of metrical structure in music, with the hierarchy ordering the levels of rhythmical values such as whole notes, half notes, quarter notes, etc.

We track large part of the metrical structure, by following in parallel each metrical level separately and combining all the levels in one single hierarchical structure. In this metrical hierarchy, a limited number of metrical levels are detected as dominant levels, for particular periods of time in the piece of music being analyzed. Dominant levels might sometimes correspond to what previous approaches consider as tactus and bar beats.

The internal model of metrical hierarchy considers that pulse lags of individual metrical levels are in exact integer relation one with the others. Apart from the first dominant metrical level discovered i_0, each metrical level i is dependent on another metrical level i_{r_i}: its theoretical pulse lag $\widehat{\tau^i}$ is at any time instant n a multiple or division of its referential metrical level:

$$\widehat{\tau^i_n} = \widehat{\tau^{i_{r_i}}_n} \times m^i \text{ or } \widehat{\tau^i} = \widehat{\tau^{i_{r_i}}_n}/d^i \quad (1)$$

The pulse lags of the entire metrical hierarchy at a time instant n is therefore conditioned solely by the pulse lag $\widehat{\tau^{i_0}_n}$ of one single level i_0, associated with the first dominant level discovered.

$$\widehat{\tau^i_n} = \widehat{\tau^{i_0}_n} \times l^i \quad (2)$$

5.1 Causal algorithm

The analysis is causal: the whole process is carried out for each successive time instant, during which all the levels of the metrical hierarchy are tentatively mapped with the peaks of the periodicity curve at that given time frame n, i.e. real lag values of the form τ^i_n are given to the different levels i. In the same time, the theoretical set of values given by equation 2 are updated so that they map as closely as possible with the real values.

For each successive time frame n, peaks k in the periodicity function are considered in decreasing order of periodicity score p_k.

Each peak k, related to a periodicity lag t_k is tentatively mapped to one metrical level i. For that aim, a succession of tests is carried out.

- The first test explained in [3] has been removed.
- We try to associate the peak to any currently active metrical level $i \in A$:

  $$i^A_k = \arg\min_{i \in A} \left(\min(|t_k - \tau^i_*|, |t_k - t_{k-1}|)\right) \quad (3)$$

  where τ^i_* indicates the current periodicity lag value associated with level i, it can be τ^i_n if there has already been a peak at the current time frame n associated with that level, or else its value at the most recent frame where a peak was found $\tau^i_{n-m}, m < M$.

  - If no peak has been integrated into the metrical hierarchy at the current time frame n, the chosen metrical level is candidate to become dominant, which would be likely only if this integration is particularly smooth:

    $$\left|t_k - \tau^{i^A_k}_*\right| < .1 \text{ and } \left|\log_2\left(\frac{t_k}{\tau^{i^A_k}_*}\right)\right| < .2 \quad (4)$$

    $$\implies \tau^{i^D_k}_n = t_k \quad (5)$$

    If this succeeds, the chosen metrical level is considered as dominant if its current peak periodicity score is sufficiently high and if its referential metrical level is also already dominant:

    $$p_k > \theta \text{ and } i_{r_{i^A_k}} \in D \implies i^A_k \in D \quad (6)$$

  - In the other cases, this integration can be considered under a loosen condition:

    $$\left|t_k - \tau^{i^A_k}_*\right| < \delta_k \implies \tau^{i^A_k}_n = t_k \quad (7)$$

    If the current peak has a pulse lag t_k that is closer to the theoretical peak than the currently registered metrical level lag $\widehat{\tau^{i^A_k}_n}$ is, then the metrical level is updated:

    $$\left|t_k - \tau^{i^A_k}_n\right| < \left|\widehat{\tau^{i^A_k}_n} - \tau^{i^A_k}_n\right| \implies \tau^{i^A_k}_n = t_k \quad (8)$$

If that same peak is sufficiently strong ($p_k > .1$), we check whether it initiates a new metrical level:

- For all the slower metrical levels, we find those that have a theoretical pulse lag that is in integer ratio with the peak lag:

  $$i \in A, \min\left(\frac{\widehat{\tau^i_n}}{t_k} \bmod 1, 1 - \left(\frac{\widehat{\tau^i_n}}{t_k} \bmod 1\right)\right) < \epsilon \quad (9)$$

  where ϵ is set to to .02 if no other stronger peak in the current time frame n has been identified with the metrical hierarchy, and else to .2 in the other case.

  If we find several of those slower levels in integer ratio, we select the fastest one, unless we find a slower one with a ratio defined in equation 9 that would be closer to 0.

- Similarly, for all the faster metrical levels, we find those that have a theoretical pulse lag that is in integer ratio with the peak lag:

  $$i \in A, \min\left(\frac{t_k}{\widehat{\tau^i_n}} \bmod 1, 1 - \left(\frac{t_k}{\widehat{\tau^i_n}} \bmod 1\right)\right) < \epsilon \quad (10)$$

  where ϵ is set to to .02 if no other stronger peak in the current time frame n has been identified with the metrical hierarchy, and else to .2 in the other case.

- If we have found both a slower and a faster level, we select the one with stronger periodicity score.

- This gives us a referential metrical level i_R, upon which our new discovered metrical level i_N will be based. The level index l_{i_N} of the new metrical level is defined as:

$$l_{i_N} = l_{i_R} * \left[\frac{t_k}{\widehat{\tau_n^{i_R}}} \right] \quad (11)$$

Finally, if the strongest periodicity peak in the given time frame n is not associated with any level of the metrical hierarchy, a new metrical hierarchy is created, with a single metrical level related to that peak. These multiple metrical hierarchies live parallel existences, and the algorithm continues by tentatively mapping the peaks of the periodicity curve on these multiple hierarchies in parallel. Mechanisms have also been conceived to fuse multiple hierarchies whenever it turns out that they belong to a single hierarchy.

Once all the peaks p_k of a given time frame n have been considered, the theoretical pulse lags are updated based on the new empirical data collected.

For lack of space, the details of the models are not given in this paper. Values used for some parameters defined in this section: $\delta_0 = .1, \theta = .15, M = 10$.

Complete examples of metrical structures are shown and discussed in [2].

In *MIRtoolbox* 1.5 and later, this metrical analysis can be performed by simply calling the new *mirmetre* operator.

6. TEMPO-RELATED METRICAL LEVEL SELECTION

If several metrical hierarchies have been constructed on a given musical except, the metrical hierarchy covering the largest temporal span is selected for the definition of the tempo.

For the selected metrical hierarchy, to each metrical level is associated a numerical score, computed as a summation across frames of the related periodicity score for each frame. An additional weight proposed in the previous submission [3] has been removed.

In the variant OL2, scores that are below a given threshold (set to 2) are filtered out.

Then we construct all possible metrical structures, made of a series of levels that have integer ratio, and that could be related to the idea of tactus/tatum/bar decomposition. We choose the metrical structure that yields the best overall score (obtained by summing the score related to each selected level).

One problem with this method is that candidate metrical hierarchies can have various number of levels, so comparing the summation of the score would penalise those with fewer number of levels. Therefore, while our prevous submission [3] was simply comparing these summation, in the new models OL1 and OL2, when comparing two hierarchies, we only select the most dominant levels of each hierarchy in such a way that we get two hierarchies with same number of levels. This truncation is virtually performed only when comparing pair of hierarchies; the select hierarchy keeps all its levels that were initially constructed.

For the selected metrical structure, we finally select two most dominant levels with periodicity between 30 and 300 BPMs by choosing the two levels with best scores. Before selection, these scores are first weighted by a resonance curve that indicates a preference for periodicities closer to 120 BPMs.

In the new models OL1 and OL2, a periodicity that is higher than 140 BPM cannot belong to the two selected metrical levels, except if that fast pulsation is ternary, i.e., if the pulsation at the next level is three times lower.

For each of the two selected metrical levels, the final tempo value is obtained by taking the median of the BPM values collected across frames.

7. REFERENCES

[1] O. Lartillot, and P. Toiviainen: "A Matlab Toolbox for Musical Feature Extraction From Audio," *Proceedings of the International Conference on Digital Audio Effects (DAFx 2007)*, 2007.

[2] Olivier Lartillot, Donato Cereghetti, Kim Eliard, Wiebke J. Trost, Marc-Andr Rappaz, and Didier Grandjean: "Estimating tempo and metrical features by tracking the whole metrical hierarchy," *Proceedings of the 3rd International Conference on Music & Emotion (ICME 2013)*, 2013.

[3] Olivier Lartillot: "Mirtempo 1.6 : Tempo Estimation by Tracking a Complete Metrical Structure Using a Rich Onset Detector," *MIREX*, 2013.

[4] O. Lartillot, *MIRtoolbox User's Manual*, http://www.jyu.fi/hum/laitokset/musiikki/en/research/coe/materials/mirtoolbox