

# Context-Dependent Beat Tracking of Musical Audio

Matthew E. P. Davies, *Student Member, IEEE*, and Mark D. Plumbley, *Member, IEEE*

**Abstract**—We present a simple and efficient method for beat tracking of musical audio. With the aim of replicating the human ability of tapping in time to music, we formulate our approach using a two state model. The first state performs tempo induction and tracks tempo changes, while the second maintains contextual continuity within a single tempo hypothesis. Beat times are recovered by passing the output of an onset detection function through adaptively weighted comb filterbank matrices to separately identify the beat period and alignment. We evaluate our beat tracker both in terms of the accuracy of estimated beat locations and computational complexity. In a direct comparison with existing algorithms, we demonstrate equivalent performance at significantly reduced computational cost.

**Index Terms**—Beat tracking, musical meter, onset detection, rhythm analysis.

## I. INTRODUCTION

THE ACT of tapping one's foot in time to music is an intuitive and often unconscious human response [1]. The computational equivalent to this behavior, commonly referred to as *beat tracking*, has the aim of recovering a sequence of beat onset times from a musical input consistent with human foot taps. This task is related both to note onset detection [2], that of identifying the start points of musical events in an audio signal, and tempo induction [3], that of finding the underlying rate of a piece of music, and forms part of research into automatic rhythm description [4].

The need for robust beat tracking extends beyond the extraction of a sequence of beat onset times. In particular, within music information retrieval (MIR)-based research, automatically extracted beats can provide a musically meaningful temporal segmentation for further analysis, such as chord estimation for harmonic description [5], long-term structural segmentation of audio [6], [7], higher level metrical analysis [8], and rhythmic genre classification [9], [10].

If we consider a black-box beat tracker which receives a musical input and produces a sequence of beat times, we can suggest several desirable properties for such a system:

- 1) both audio and symbolic musical signals can be processed;
- 2) no *a priori* knowledge of the input is required (e.g., regarding genre, timbre, or polyphony);

Manuscript received March 30, 2006; revised July 5, 2006. the work of M. Davies was supported by a college studentship from Queen Mary University of London, U.K. This work was supported in part by the EU-FP6-IST-507142 under Project SIMAC and in part by EPSRC under Grants GR/S75802/01 and GR/S82213/01. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Rudolf Rabenstein.

The authors are with the Department of Electronic Engineering, Queen Mary, University of London, London E1 4NS, U.K. (e-mail: matthew.davies@elec.qmul.ac.uk; mark.plumbley@elec.qmul.ac.uk).

Digital Object Identifier 10.1109/TASL.2006.885257

- 3) perceptually accurate beat locations can be identified, efficiently and in real-time if necessary;
- 4) changes in tempo can be followed, such as the result of step or ramp changes and expressive timing variation.

Despite a large body of research into beat tracking (e.g., [8], [11]–[13]), a system that meets all of these requirements does not currently exist.

## A. Prior Art

Early approaches to beat tracking and rhythm analysis (e.g., [11], [14], and [15]) processed symbolic data rather than audio signals, perhaps due to limited computational resources or a lack of sophistication in note onset detection. Recent advances in onset detection and audio feature extraction (reviewed in [2]) have enabled the use of audio signals for beat tracking. We now present a brief review of five well-known approaches to beat tracking of audio signals. For more details, see [4].

Dixon's approach, *Beatroot* [1], processes a sequence of note onset times either extracted from an audio signal or from a symbolic representation, within a multi-agent system. Likely tempo hypotheses are derived from clustering inter-onset-intervals. These are used to form multiple beat agents (with varying tempo and phase) which compete based on how well each can predict beat locations. Dixon's algorithm is designed to track beats in expressively performed music.

Goto's approach [8] is also agent based. In addition to tracking the beats (at the 1/4 note level) analysis is extended to the 1/2 and whole note levels. Onset analysis is performed across seven parallel subbands, where spectral models are used to extract snare and bass drum events. Chord changes and predefined rhythmic pattern templates are used to infer the beats and higher level metrical structure. Goto's system operates accurately and in real-time, provided the input signal has a steady tempo and a 4/4 time-signature (i.e., four beats per bar).

Hainsworth [16] extracts note onsets from subband analysis using two distinct models, one for finding transient events, the other designed to detect harmonic change, which are combined to give a single sequence of onsets. Particle filters are used within a statistical framework to track the beats, which are modelled as a quasi-periodic sequence driven by time-varying tempo and phase processes. The main limitation of this approach is its computational complexity.

In contrast to the previous approaches, which rely on the automatic extraction of note onsets, Scheirer [12] proposes the use of a psychoacoustically motivated amplitude envelope signal to emphasise (but not extract) note onset locations, as the front-end to his beat tracker. Six octave-spaced subband envelope signals are passed through a parallel bank of tuned comb filter resonators representing tempo hypotheses over the range of 60–180

bpm. These phase-lock to the envelope signals across the subbands, to simultaneously infer the tempo and predict the phase of the beats within a real-time system.

Klapuri *et al.* [13] expand upon Scheirer's amplitude envelope and comb filter model. They adopt a more robust *registral accent signal* across four parallel analysis bands as the input to their system and use comb filterbanks within a probabilistic framework to simultaneously track three metrical levels. These correspond to the *tatum*, the lowest or fastest metrical level, the *tactus*, the rate at which humans are most likely to tap, and the *measure*, which indicates the grouping of beats into bars. Analysis can be performed causally and noncausally, and is not restricted to any particular genre, tempo or time-signature. The algorithm has been tested over a large annotated database with results demonstrating superior performance to the Scheirer [12] and Dixon [1] approaches, while Hainsworth [16] reports comparable accuracy (a comparison against the Goto system [8] is not included).

A comprehensive comparison of audio beat trackers has not yet been undertaken. However, in a recent study of audio tempo induction algorithms [3], many of which inferred the tempo by first extracting the beats, the approach of Klapuri *et al.* [13] was shown to be most accurate. We therefore identify this method as representative of the current state of the art for beat tracking of audio signals.

## B. Proposed Model

A key component in audio based beat trackers appears to be the use of a multiband decomposition prior to any rhythmic analysis. Scheirer [12] proposed that the use of subbands was essential; that a single analysis band was not sufficient to preserve the rhythmic content when using an amplitude envelope-based input representation. We demonstrate that by using a more sophisticated input representation, derived across a wide analysis band, we can provide a single input to our beat tracker, without the need for subband analysis. We use the complex spectral difference onset detection function [17], described in Section II-A.

To extract the beats from this input representation, we could attempt to recover the beat period and phase simultaneously, as in [12] and [13] or separate this task into two individual processes [1], [8]. The latter is computationally more efficient, as it decomposes a two-dimensional search into two, one-dimensional searches. We estimate the beat period by passing the autocorrelation function of the onset detection function through a shift invariant comb filterbank matrix, as described in Section II-B1. We then use the beat period to recover the beat alignment by passing the onset detection function through a comb filter matrix described in Section II-B2.

Klapuri *et al.* [13] show that the robustness of their meter analysis model is due to the probabilistic modelling of the temporal evolution and interaction between each of the three metrical levels analysed. To increase the reliability of the beat output within our beat tracker, we adopt a simpler, heuristic approach by embedding context-dependent information directly into the beat period and alignment estimation processes. This is described in Sections II-C and II-D. An overview of our approach is shown in Fig. 1.

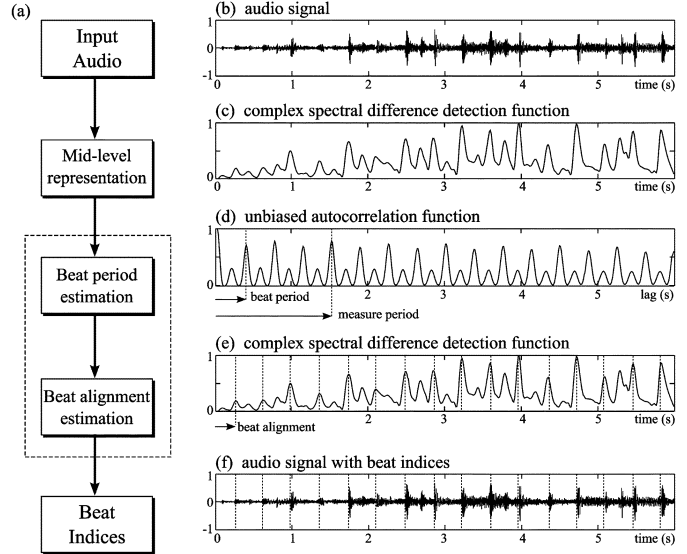


Fig. 1. (a) Flow chart of typical beat tracking process. (b)-(f): Proposed model. (b) Input audio signal. (c) Onset detection function. (d) Autocorrelation function with beat period. (e) Detection function with beat alignment. (f) Audio signal with dotted vertical lines showing beat locations.

We evaluate our proposed beat tracking system in Section III, both in terms of the accuracy of generated beats and computational complexity. We demonstrate equivalent performance to the current state of the art, but at a significant reduction in computational cost, presenting discussions and conclusions in Sections IV and V.

## II. APPROACH

We recover beat locations from a continuous midlevel representation derived from a musical audio input, described in Section II-A. Our approach to beat tracking is comprised of two states. The first of these, which we refer to as the *General State*, performs beat period and beat alignment induction without any prior knowledge of the input (see Section II-B). We then describe the incorporation of context-dependent information regarding the tempo, time-signature and the locations of past beats, into the extraction of these parameters as part of the *Context-Dependent State* (Section II-C). The higher level operation of the beat tracking system, which controls the interaction between the General and Context-Dependent States is given in Section II-D.

### A. Midlevel Input Representation

A common first stage in audio beat tracking algorithms is the transformation of the input audio signal into a more meaningful, compact representation from which to infer the beats. This can take the form of a sequence of automatically extracted note onsets [16], or a continuous function that emphasizes them by exhibiting local maxima at likely onset locations [12]. These continuous signals, often referred to as *onset detection functions* [2], can more generally be considered *midlevel representations* that act as an intermediary signal between the input audio and the output beats. The process of extracting onsets from an onset

detection function is not 100% effective, and can introduce undesirable errors both in terms of missed detections and false positives. We therefore use a continuous detection function as the input to our beat tracker. To allow for the widest range of input signals, we would like the onset detection function to be sensitive both to percussive events (e.g., a drum hit) and softer tonal onsets (e.g., from a bowed violin). Referring to a recent comparative study of onset detection algorithms [2], we select the *complex spectral difference* onset detection function as that which best matches our aims. The ability of this detection function to emphasise note onsets relies on the extent to which the spectral properties of the signal at the onset of musical events are unpredictable. By working in the complex domain, note onsets are emphasized either as a result of significant change in energy in the magnitude spectrum, and/or a deviation from the expected phase values in the phase spectrum, caused by a change in pitch. The detection function  $\Gamma(m)$ , at frame  $m$ , is calculated by measuring the sum of the Euclidean distance between the observed spectrum at frame  $S_k(m)$  and a prediction of the spectrum  $\hat{S}_k(m)$ , for all frequency bins  $k$

$$\Gamma(m) = \sum_{k=1}^K |S_k(m) - \hat{S}_k(m)|^2. \quad (1)$$

To allow the beat tracker to operate independently of the sampling frequency of the input, the detection function is derived with a fixed time resolution. We set this resolution equal to that used in recent work in onset detection [2] at  $t_{\text{DF}} = 11.6$  ms per detection function (DF) sample. A derivation of the detection function is given in Appendix I, with further details in [17]. An example audio signal and detection function are shown in Fig. 1(b) and (c).

To permit the tempo and phase of the beats to vary across the length of the input, we partition the detection function  $\Gamma(m)$  into overlapping analysis frames  $\Gamma_i(m)$

$$\Gamma_i(m) = \begin{cases} \Gamma(m), & m = 1 + (i-1)B_f \dots, B_f + (i-1)B_h \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where frames are of length  $B_f = 512$  DF samples with hop size  $B_h = B_f/4$  giving an overlap of 75%. Given the fixed time-resolution of the onset detection function, this equates to a 6-s analysis frame with a 1.5-s step increment, where each analysis frame is long enough to reliably estimate the tempo and the step increments short enough to track any changes that might occur.

### B. General State

Within our proposed beat tracking model we use context-dependent information regarding the tempo and past beat locations to improve the reliability of the extracted beat locations. However, in order to incorporate this information into the beat tracking model, we must first extract it from the input signal. For this purpose, we propose a *General State* for beat tracking. Given an input onset detection function frame, the role of the General State is to infer the time between successive beats, the *beat period*, and the offset between the start of the frame and

the locations of the beats, the *beat alignment*, without any prior knowledge of the input.

1) *Beat Period Induction*: Other approaches to beat tracking have simultaneously extracted the beat period and beat alignment from a midlevel representation using comb filter resonators [12], [13]. We adopt a more efficient approach by first identifying the beat period and then subsequently finding the beat alignment. To facilitate this separation, we discard all phase-related information using an autocorrelation function of the current detection function frame. To achieve an autocorrelation function with clear, well-defined peaks, we first process the detection function to emphasise the strongest and discard the least significant peaks. We calculate an adaptive, moving mean threshold  $\bar{\Gamma}_i(m)$

$$\bar{\Gamma}_i(m) = \text{mean}\{\Gamma_i(q)\} \quad m - \frac{Q}{2} \leq q \leq m + \frac{Q}{2} \quad (3)$$

where  $Q = 16$  DF samples. We then subtract  $\bar{\Gamma}_i(m)$  from the original detection function and half-wave rectify the output, setting any negative valued elements to zero, to give a modified detection function frame  $\tilde{\Gamma}_i(m)$

$$\tilde{\Gamma}_i(m) = \text{HWR}(\Gamma_i(m) - \bar{\Gamma}_i(m)) \quad (4)$$

where  $\text{HWR}(x) = (x + |x|)/2$ .

We calculate an autocorrelation function  $A(l)$ , normalized to correct the bias occurring as a linear function of increasing lag

$$A(l) = \frac{\sum_{m=1}^{B_f} \tilde{\Gamma}_i(m) \tilde{\Gamma}_i(m-l)}{|l - B_f|} \quad l = 1, \dots, B_f. \quad (5)$$

We can map a lag  $l$  (in DF samples) from the autocorrelation domain into a measure of tempo, in beats per minute, using the following relationship:

$$\text{bpm} = \frac{60}{l \times t_{\text{DF}}} \quad (6)$$

where  $t_{\text{DF}} = 11.6$  ms is the resolution of the detection function, specified in Section II-A.

Brown [18] uses autocorrelation to infer the meter of musical scores and symbolic performances observing periodicity at integer multiples of the beat level with a strong peak at the measure or bar level. We adopt a similar approach to finding the beat period, using comb filter-type structures. To reflect periodicity at several metrical levels, we derive a comb template  $\lambda_\tau(l)$  as the sum of weighted delta functions at integer multiples of an underlying periodicity  $\tau$

$$\lambda_\tau(l) = \sum_{p=1}^4 \sum_{v=1-p}^{p-1} \frac{\delta(l - \tau p + v)}{2^p - 1} \quad l = 1, \dots, B_f. \quad (7)$$

We define the longest observable periodicity  $\tau_{\text{max}} = B_h$  DF samples, which is equivalent to 40 bpm from (6). Equating  $\tau_{\text{max}}$  to the analysis frame hop size means, we will derive at least one beat from every analysis frame. Given each autocorrelation function frame is of length  $B_f = 512$  DF samples and  $\tau_{\text{max}} = B_h$ , this permits the use of four comb elements  $p = 1, \dots, 4$  within each comb template  $\lambda_\tau(l)$ . To account for




Fig. 2. Autocorrelation function with metrically unbiased comb template.

the poor resolution of the autocorrelation function at short lags, each comb element has a width proportional to its relationship to  $\tau$  set by  $v = 1 - p, \dots, p - 1$ , and a height normalized by its width  $2p - 1$ . We refer to this comb template as *metrically unbiased* because each comb element has equal influence over the extracted beat period. An example comb template and autocorrelation function are given in Fig. 2.

We combine a number of comb templates to span the range of beat period hypotheses ( $1 \leq \tau \leq B_h$ ) into a *shift-invariant* comb filterbank  $F_G(l, \tau)$ . To reflect an approximate prior distribution of beat period hypotheses, we use a weighting curve  $w_G(\tau)$  to indicate those beat periods which are most likely. Other systems have used a log-Gaussian distribution [13] or a resonance function [19] for this purpose. We require a function which strongly attenuates short lags while slowly decaying over longer lags and formulate  $w_G(\tau)$  from the Rayleigh distribution function

$$w_G(\tau) = \frac{\tau}{\beta^2} e^{-\tau^2/2\beta^2} \quad \tau = 1, \dots, B_h \quad (8)$$

where the  $\beta$  parameter sets the strongest point of the weighting. An acceptable range for  $\beta$  was found to be between 40 and 50 DF samples, from which we select  $\beta = 43$ , equivalent to 120 bpm from (6). Informal tests indicated the resonance curve [19] and log-Gaussian distribution [13] to be less effective than the Rayleigh weighting within our beat tracking framework.

The shift-invariant comb filterbank  $F_G(l, \tau)$  can be conveniently represented in matrix form, where each column of the matrix represents a comb template  $\lambda_\tau(l)$  with a fundamental periodicity  $\tau$  equal to the column number. We apply the weighting curve multiplying each  $\lambda_\tau(l)$  with the appropriate value from  $w_G(\tau)$ , to give

$$F_G(l, \tau) = w_G(\tau) \lambda_\tau(l). \quad (9)$$

We take the product of the autocorrelation function  $A(l)$  with  $F_G(l, \tau)$  to give an output  $y_G(\tau)$  as a function of  $\tau$

$$y_G(\tau) = \sum_{l=1}^{B_f} A(l) F_G(l, \tau) \quad (10)$$

from which we identify the beat period  $\tau_G$  as the index of the maximum value

$$\tau_G = \arg \max_{\tau} (y_G(\tau)). \quad (11)$$

The comb filterbank  $F_G(l, \tau)$  is shown in Fig. 3(a), with an example output signal  $y_G(\tau)$  in Fig. 3(b). We can observe multiple peaks in  $y_G(\tau)$ , which correspond to integer and fractional

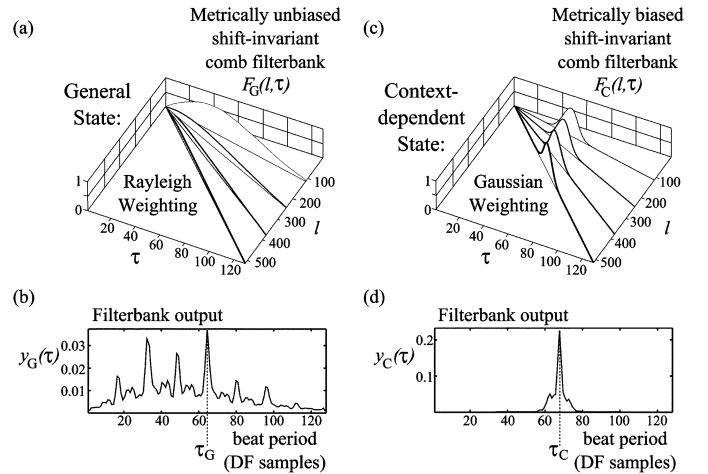


Fig. 3. Shift-invariant comb filterbanks for the (a) General State and (c) the Context-Dependent State, with respective output signals (b) and (d). The index of the maximum value in the output is selected as the beat period for the current analysis frame.

multiples of  $\tau_G$ . The General State relies on the lag consistent with the tempo of the input being the strongest peak within  $y_G(\tau)$  for all analysis frames.

2) *Beat Alignment Induction*: We define beats in terms of two parameters: the beat period  $\tau_G$ , and the beat alignment  $\alpha_G$ . We now address finding the locations of the beats within each frame given a known beat period. In our beat tracking system, we can perform either causal or noncausal analysis. To distinguish between causal and noncausal analysis, we use  $x^*$  to indicate parameter  $x$  is extracted causally.

For noncausal analysis, each frame contains  $B_f$  DF samples of future data. We define the beat alignment  $\alpha_G$  to represent the offset from the start of the current frame to the location of the first beat within this frame, as shown in Fig. 1(e). For causal analysis, each frame contains the past  $B_f$  DF samples of data. Without access to any future data, we predict the locations of future beats based solely on analysis of past data. We define the causal beat alignment  $\alpha_G^*$  as the offset from the end of current frame back to the last beat occurring in this frame. This represents the best location from which future beats can be predicted.

To extract the beat alignment, we adopt a similar matrix-based approach to finding the beat period. For beat period extraction, each column of the matrix  $F_G(l, \tau)$  represents a comb template with a different periodicity. Once this periodicity is known, we can formulate a comb filter matrix  $H_G(m, \alpha)$  containing all possible shifts of the beat period to infer the beat alignment from the current detection function frame  $\Gamma_i(m)$ .

We construct an alignment comb template  $\psi_\alpha(m)$  as a sequence of impulses at beat period intervals with offset  $\alpha$ . Since we cannot guarantee that the tempo will be constant across the entire frame, we multiply each comb template with a linearly decreasing weighting  $\nu(m)$  to give most emphasis to the first few comb elements

$$\psi_\alpha(m) = \sum_{k=1}^{\lfloor B_f/\tau_G \rfloor} \nu(m) \delta(m - k\tau_G + \alpha) \quad m = 1, \dots, B_f \quad (12)$$

where  $k$  ranges over the number of elements in each comb template and  $\nu(m)$  is defined as

$$\nu(m) = \frac{B_f - m}{B_f} \quad m = 1, \dots, B_f. \quad (13)$$

To form the matrix  $H_{\mathbf{G}}(m, \alpha)$  each column represents a comb template  $\psi_{\alpha}(m)$  with an incremental offset  $\alpha$

$$H_{\mathbf{G}}(m, \alpha) = \psi_{\alpha}(m) \quad 1 \leq \alpha \leq \tau_{\mathbf{G}}. \quad (14)$$

For noncausal analysis, we find the product of the current detection function frame  $\Gamma_i(m)$  with  $H_{\mathbf{G}}(m, \alpha)$ , to give an output function of  $\alpha$

$$z_{\mathbf{G}}(\alpha) = \sum_{m=1}^{B_f} \Gamma_i(m) H_{\mathbf{G}}(m, \alpha) \quad (15)$$

with  $\alpha_{\mathbf{G}}$  extracted as the index of the maximum value of  $z_{\mathbf{G}}(\alpha)$

$$\alpha_{\mathbf{G}} = \arg \max_{\alpha} (z_{\mathbf{G}}(\alpha)). \quad (16)$$

We then place beats  $\gamma_{i,b}$  for the  $i$ th frame at beat period intervals from  $\alpha_{\mathbf{G}}$  such that

$$\gamma_{i,b} = (t_i + \alpha_{\mathbf{G}}) + (b - 1)\tau_{\mathbf{G}} \quad b = 1, \dots, B \quad (17)$$

where  $t_i$  indicates the time at the start of the  $i$ th detection function frame,  $\Gamma_i(m)$ . To allow for variations in the beat structure, any beats beyond one step increment from  $t_i$  will be assigned from analysis of the subsequent detection function frame  $\Gamma_{i+1}(m)$ , we therefore constrain  $\gamma_{i,B} < t_i + B_h$ .

For the causal version, we modify (15), replacing the current detection function frame  $\Gamma_i(m)$  with a time-reversed version  $\Gamma_i^*(m) = \Gamma_i(-m)$  leaving  $H_{\mathbf{G}}(m, \alpha)$  unchanged. This now gives strongest emphasis to the comb elements at the end of each frame, which are the most recent DF samples to have occurred. We find  $z_{\mathbf{G}}^*(\alpha)$ , the causal version of  $z_{\mathbf{G}}(\alpha)$  as

$$z_{\mathbf{G}}^*(\alpha) = \sum_{m=1}^{B_f} \Gamma_i^*(m) H_{\mathbf{G}}(m, \alpha) \quad (18)$$

extracting  $\alpha_{\mathbf{G}}^*$  as the index of the maximum value

$$\alpha_{\mathbf{G}}^* = \arg \max_{\alpha} (z_{\mathbf{G}}^*(\alpha)). \quad (19)$$

We modify (17) to predict causal beats  $\gamma_{i,b}^*$  up to one-step increment ( $B_h$  DF samples) beyond the end of the current frame

$$\gamma_{i,b}^* = (t_i^* + \alpha_{\mathbf{G}}^*) + b\tau_{\mathbf{G}} \quad b = 1, \dots, B \quad (20)$$

where  $\gamma_{i,B}^* < t_i^* + B_h$  and  $t_i^*$  represents the time at the end of the frame  $i$ .

The comb filter matrix  $H_{\mathbf{G}}(m, \alpha)$  is shown in Fig. 4(a) with a corresponding output signal  $z_{\mathbf{G}}(\alpha)$  in Fig. 4(b). We can observe two peaks in the output signal, which are labeled as the *on-beat*, the result of tapping in phase, and the *off-beat*, tapping in antiphase, i.e., at an offset of  $\tau_{\mathbf{G}}/2$  DF samples from

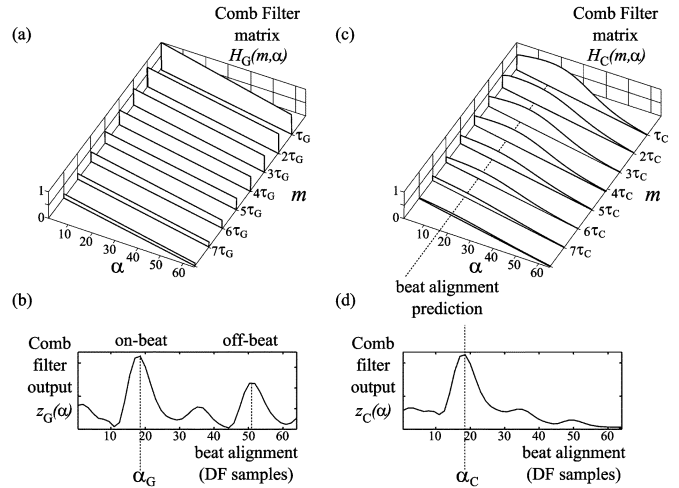


Fig. 4. Comb filter matrices for the (a) General State and (c) Context-Dependent State, with respective example output signals (b) and (d). The index of the maximum value in the output is selected as the phase of the beats for the current analysis frame.

$\alpha_{\mathbf{G}}$  ( $\tau_{\mathbf{G}}/2 \approx 30$  here). We assume that the on-beat will be the strongest peak in  $z_{\mathbf{G}}(\alpha)$  for all analysis frames. To maintain a consistent beat output it is important not to erroneously switch between the on and the off-beat.

### C. Context-Dependent State

The General State for beat tracking operates in a memoryless fashion, extracting the beat period and beat alignment through a process of repeated induction. The principal limitation of this method is that no effort is made to enforce continuity. For a given analysis frame  $\tau_{\mathbf{G}}$  and  $\alpha_{\mathbf{G}}$  can vary independently of any past values. This can cause two common beat tracking errors [20]: 1) the switching of metrical levels, i.e., tapping at one rate and then changing to half or double this rate without a tempo change having occurred and 2) switching between the on and the off-beat, or in other words changing from tapping in phase to tapping beats in anti-phase.

To prevent these undesirable effects we propose a *Context-Dependent State* for beat tracking. It operates in a similar way to the General State in that it separately extracts the beat period and the beat alignment. We now incorporate prior knowledge of the tempo of the input and feedback past locations of beats, based on the simple assumption that, if beats are arriving at regular beat period intervals, then our focus is no longer on inducing beat locations but in maintaining continuity [21].

1) *Beat Period*: There are two principle distinctions between the process of beat period extraction in the Context-Dependent State and the General State both of which relate to the construction of the context-dependent comb filterbank  $F_C(l, \tau)$ . All other aspects involved in the extraction of the beat period from the autocorrelation function  $A(l)$  remain constant. First, we classify the time-signature of the input and use this to incorporate metrical bias into each comb template, setting the number of elements equal to the numerator of the time-signature. Second, we replace the Rayleigh weighting curve, which gave an approximate prior distribution of beat periods, with a

tighter Gaussian weighting localized around a known beat period. This prevents the extraction of beat periods at related metrical levels.

To extract the time-signature  $T$  of the input we use a similar approach to the method of Brown [18]. This method extracts a strong peak in the autocorrelation function at the measure level periodicity and relates this to the beat period. We follow Gouyon and Herrera [22] and classify the input as having a *binary* meter  $T = 4$ , or a *ternary* meter  $T = 3$ . We do not attempt to infer odd time-signatures (e.g.,  $5/4$ ,  $7/8$ ,  $9/8$ ), leaving this as a topic for further work.

Given a value for the beat period either from analysis in the General State  $\tau_G$  or from a notated value, we classify  $T$  by comparing the energy at odd and even multiples of the beat period within the autocorrelation function based on the following condition:

$$A(2\tau_G) + A(4\tau_G) > A(3\tau_G) + A(6\tau_G). \quad (21)$$

If condition (21) holds, i.e., we find stronger periodicity at even multiples than odd multiples, we infer a binary meter setting  $T = 4$ ; otherwise, we presume a ternary meter and set  $T = 3$ . This allows us to reformulate the comb filterbank matrix  $F_C(l, \tau)$ . We set the number of elements in each comb template  $\lambda_\tau(l)$  equal to  $T$ . To make the comb template metrically biased [in contrast to (7)] we remove the  $(2p - 1)$  normalization from (7) to give most emphasis to the measure level periodicity

$$\lambda_\tau(l) = \sum_{p=1}^T \sum_{v=1-p}^{p-1} \delta(l - \tau p + v) \quad l = 1, \dots, B_f. \quad (22)$$

Prior knowledge of the beat period allows us to limit the range of likely beat period hypotheses replacing the Rayleigh weighting with a context-dependent Gaussian function centered on  $\tau_G$

$$w_C(\tau) = e^{-(\tau - \tau_G)^2 / 2\sigma_w^2} \quad \tau = 1, \dots, B_h. \quad (23)$$

The width of the weighting, defined by  $\sigma_w$ , constrains the range of likely beat periods. If  $\sigma_w$  is too tight, no local variation in the beat period will be observable. If it is too wide, the weighting will not be able to prevent the beat period switching between metrical levels. To give an approximate measure of local variation in beat periods, we analyzed the distribution of inter-annotation-intervals,  $\Delta_j$ , (normalized to have zero mean) across the annotated test database we use to evaluate our beat tracker (in Section III). We set  $\sigma_w = 4$  DF samples equating it to the standard deviation of the inter-annotation-interval distribution. Informal tests revealed  $\sigma_w$  could vary by a factor of 2 without altering the overall performance of the beat tracker.

The shift-invariant comb filterbank  $F_C(l, \tau)$  for the Context-Dependent State can now be formed by setting each comb template  $\lambda_\tau(l)$  as the  $\tau$ th column of the matrix and weighting each template using the Gaussian weighting  $w_C(\tau)$

$$F_C(l, \tau) = w_C(\tau)\lambda_\tau(l). \quad (24)$$

As in the General State, we take the product of the autocorrelation function  $A(l)$  with  $F_C(l, \tau)$

$$y_C(\tau) = \sum_{l=1}^{B_f} A(l)F_C(l, \tau) \quad (25)$$

extracting  $\tau_C$  as the index of the maximum value of  $y_C(\tau)$

$$\tau_C = \arg \max_{\tau} (y_C(\tau)). \quad (26)$$

We may now compare the comb filterbanks for the General and Context-Dependent States by examining Fig. 3(a) and (c). The effect of the Gaussian weighting is evident by comparing the respective example comb output signals from each state in Fig. 3(b) and (d). For the General State, there is a strong ambiguity between two beat period hypotheses, with either one a possible interpretation of the input. However, in the Context-Dependent State, only a small range around the stronger of the two remains unattenuated by the Gaussian weighting function. The restriction in likely beat periods prevents the possibility of switching between metrical levels that could occur using the General State alone.

2) *Beat Alignment*: The extraction of the beat alignment in the Context-Dependent State does not differ greatly from the General State. Again, we rely on a previously identified beat period  $\tau_C$  and use a matrix-based approach to find the alignment  $\alpha_C$  from the current detection function frame  $\Gamma_i(m)$ . The role of the Context-Dependent State is to maintain a consistent beat output. For beat period extraction, we limit the range of likely beat periods to be close to an expected value. When this is the case, we can reasonably assume that beats will occur approximately at beat period intervals. For a given frame  $\Gamma_i(m)$ , we can use this to predict a likely value for  $\alpha_C$  given the beat output from the previous frame  $\Gamma_{i-1}(m)$ . As in the General State, we can perform both noncausal and causal analysis.

In the noncausal case, the beat alignment represents the offset from the start of the current detection function frame  $\Gamma_i(m)$  to the location of the first beat within that frame. We can predict  $\alpha_C$  as approximately  $\tau_C$  DF samples beyond the last beat  $\gamma_{i-1,B}$  from the previous detection function frame  $\Gamma_{i-1}(m)$

$$\alpha_C \approx \gamma_{i-1,B} + \tau_C. \quad (27)$$

Equating (27) would cause the beats to occur precisely at beat period intervals. To allow for the location of the beats to vary, we create an alignment weighting function  $\rho_C(\alpha)$  centered on the most likely  $\alpha_C$

$$\rho_C(\alpha) = e^{-(\alpha - (\gamma_{i-1,B} + \tau_C))^2 / 2\sigma_\rho^2} \quad \alpha = 1, \dots, \tau_C. \quad (28)$$

As with the Gaussian weighting for beat period extraction (23),  $\sigma_\rho$  defines the width of the weighting. To prevent the possibility of extracting the alignment corresponding to the off-beat, which should occur approximately  $\tau_C/2$  samples away from the on-beat [see Fig. 4(b)], we use a beat period dependent width setting  $\sigma_\rho = \tau_C/4$  DF samples.

We formulate a context-dependent comb filter matrix  $H_C(m, \alpha)$ , where each column of the matrix represents an alignment comb  $\psi_\alpha(m)$  at an offset  $\alpha$

$$\psi_\alpha(m) = \sum_{k=1}^{\lfloor B_f/\tau_C \rfloor} \nu(m) \delta(m - k\tau_C + \alpha) \quad m = 1, \dots, B_f \quad (29)$$

with  $\nu(m)$  from (13) as the linear weighting function, to give

$$H_C(m, \alpha) = \rho_C(\alpha) \psi_\alpha(m). \quad (30)$$

We then take the product of  $\Gamma_i(m)$  with  $H_C(m, \alpha)$  to give an output function of  $\alpha$

$$z_C(\alpha) = \sum_{m=1}^{B_f} \Gamma_i(m) H_C(m, \alpha) \quad (31)$$

with  $\alpha_C$  found as the index of the maximum value of  $z_C(\alpha)$

$$\alpha_C = \arg \max_{\alpha} (z_C(\alpha)). \quad (32)$$

Beats  $\gamma_{i,b}$  can then be placed at intervals of  $\tau_C$  up to one-step increment into the current analysis frame using

$$\gamma_{i,b} = (t_i + \alpha_C) + (b-1)\tau_C \quad b = 1, \dots, B \quad (33)$$

with  $\gamma_{i,B} < t_i + B_h$  and  $t_i$  as the start of the current frame.

For causal analysis, where each frame represents the past  $B_f$  DF samples of data, the beat alignment  $\alpha_C^*$  indicates the offset back from the end of the current frame. If the tempo of the input is constant, then  $\alpha_C^*$  will equal the location of the last beat  $\gamma_{i-1,B}^*$  within the previous analysis frame  $\Gamma_{i-1}^*(m)$ . Allowing for timing variations, we reformulate (28) as

$$\rho_C^*(\alpha) = e^{-(\alpha - (\gamma_{i-1,B}^*))^2 / 2(\sigma_\rho^*)^2} \quad \alpha = 1, \dots, \tau_C. \quad (34)$$

When predicting the locations of future beats, we cannot directly observe expressive timing changes as in the noncausal approach. We therefore restrict the likely beat alignment values more so than in (28) by setting  $\sigma_\rho^* = \tau_C/8$  DF samples.

We create the causal comb matrix  $H_C^*(m, \alpha)$  from (30) now using  $\rho_C^*(\alpha)$  to weight the alignment comb templates  $\psi_\alpha(m)$

$$H_C^*(m, \alpha) = \rho_C^*(\alpha) \psi_\alpha(m). \quad (35)$$

We find the product of the time-reversed detection function frame  $\Gamma_i^*(m)$  with  $H_C^*(m, \alpha)$

$$z_C^*(\alpha) = \sum_{m=1}^{B_f} \Gamma_i^*(m) H_C^*(m, \alpha) \quad (36)$$

extracting  $\alpha_C^*$  as the index of the maximum value of  $z_C^*(\alpha)$

$$\alpha_C^* = \arg \max_{\alpha} (z_C^*(\alpha)). \quad (37)$$

We can now predict the locations of future beats

$$\gamma_{i,b}^* = (t_i^* - \alpha_C^*) + b\tau_C \quad b = 1, \dots, B \quad (38)$$

where  $t_i^*$  represents the end of the current analysis frame and beats are constrained  $\gamma_{i,B}^* < t_i^* + B_h$ .

Fig. 4 shows both the General and Context-Dependent States comb filter matrices (a) and (c), with corresponding output sig-

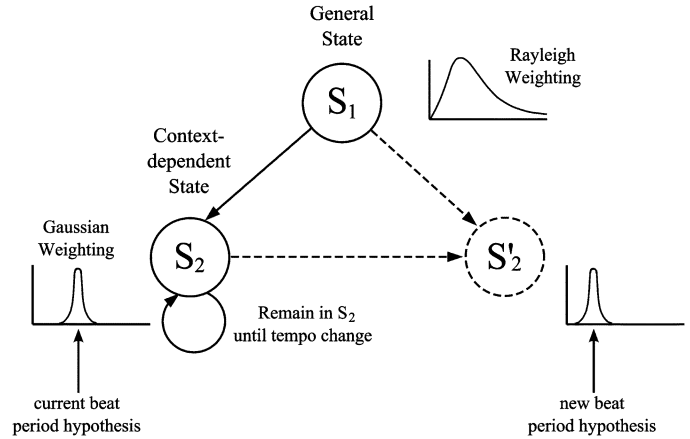


Fig. 5. Two-State switching model: Analysis begins in General State  $S_1$ . The observation of a consistent beat period generates a Context-Dependent State  $S_2$  that is replaced by  $S'_2$  if a tempo change is observed.

nals (b) and (d). The alignment weighting in the Context-Dependent State attenuates the energy in the output around the location of the off-beat leaving just a single peak to indicate the beat alignment. This prevents the possibility of the beats switching from the on-beat to the off-beat, maintaining a consistent beat output.

#### D. Two-State Model

Although the General State and Context-Dependent States can work in isolation to provide the locations of the beats, neither one is ideally suited to do so. Without any means to enforce continuity, the General State is too susceptible to switching metrical levels and changing from the on-beat to the off-beat. The Context-Dependent State, while explicitly designed to prevent these errors by limiting the range of likely beat periods and alignments, is unable to react to any tempo changes that might occur. To exploit the particular properties of each state, we combine them into a *Two-State Model*. We use the General State to infer an initial beat period and detect tempo changes, and the Context-Dependent State to maintain continuity within a given beat period hypothesis. The operation of the Two-State Model is shown in Fig. 5.

We define a state variable  $q$  to indicate which state is currently active, and therefore which is used to derive the beat locations. We label the General State  $S_1$  and the Context-Dependent State  $S_2$ . Without any prior knowledge of the input we initially set  $q = S_1$ , with beat locations assigned using the General State.

The activation of the first Context-Dependent State  $S_2$  is the result of the observation of consistent beat periods across three consecutive frames, which satisfy

$$|2\tau_G(i) - \tau_G(i-1) - \tau_G(i-2)| < \eta_1 \quad (39)$$

where  $\tau_G(i)$  is the beat period from the  $i$ th frame and  $\eta_1$  is the threshold. We define a consistent beat period to be one that can be reliably observed within the range of the Gaussian weighting (23) for beat period extraction. We therefore set  $\eta_1 = \sigma_w$  from (23) at 4 DF samples. If condition (39) is met,  $q = S_2$ , with beats assigned from the Context-Dependent State.

When  $S_2$  is active, the General State continues to operate, acting in a supervisor role, to verify whether the Context-Dependent State is appropriate to the current properties of the signal. At each frame, we extract  $\tau_G$  and compare it to  $\tau_C$ . If the values are close, then we maintain the operation of the Context-Dependent State. For the Context-Dependent State to be considered a poor match to the input, we require the following two conditions to be met: 1) the absolute difference between  $\tau_G$  and  $\tau_C$  is greater than  $\eta_2$ , the threshold for tempo changes

$$|\tau_G(i) - \tau_C(i)| > \eta_2 \quad (40)$$

where  $\eta_2$  is set to be outside of the range of the Gaussian weighting in (23), e.g.,  $\eta_2 = 2\eta_1$ ; and 2) the new beat period hypothesis  $\tau_G$  is consistently observed, and therefore meets (39).

If  $S_2$  is found to be no longer a valid match to the rhythmic properties of the signal, then it is discarded, to be replaced by a new Context-Dependent State  $S'_2$  indicative of the new beat period hypothesis observed by the General State. Before setting  $q = S'_2$  and using the new Context-Dependent State to find the beat locations, we let the General State resume control of the beat tracker for a single analysis frame, setting  $q = S_1$ , to provide a new reference point  $\alpha_G$  for the locations of the beats.

For subsequent frames, we now activate the new Context-Dependent State setting  $q = S'_2$  and use it to maintain continuity in the beat output, until a new consistent beat period is observed. Operating in this manner, the Two-State Model can explicitly model tempo discontinuities while smoothing out odd beat errors during each consistent beat period hypothesis.

### III. EVALUATION

We pursue an objective approach to evaluate our beat tracking model. We use an existing annotated test database [16] containing 222 musical audio files, with the following genre breakdown: Rock/pop (68), Dance (40), Jazz (40), Classical (30), Folk (22), and Choral (22). Each file is between 30 and 60 s in length, mono, and sampled 44.1 kHz with 16-bit resolution. The files were annotated by a trained musician, by recording taps in time to the audio files. The annotations were then synthesized over each example and where necessary, manually corrected to be perceptually *in time* according to the judgement of the annotator. For more details, see [16].

We compared the performance of our proposed model to two published beat tracking systems: 1) Simon Dixon's algorithm *Beatroot* [1] and 2) the more recent approach of Klapuri *et al.* [13], which we will refer to as KEA. To permit an informal comparison between computational and human approaches to beat tracking, we also evaluate an example human tapping performance (that of the first author). For each file in the database, beats were recorded as real-time keyboard taps using a published annotation tool [23]. In contrast to the ground truth annotated data, no manual correction of the tapped beats was permitted.

#### A. Beat Accuracy

To quantify the performance of our beat tracker, we use an objective measure to describe beat accuracy. We follow the recent approaches of [13], [16], [24] by placing an emphasis on the generation of a continuously correct beat output.

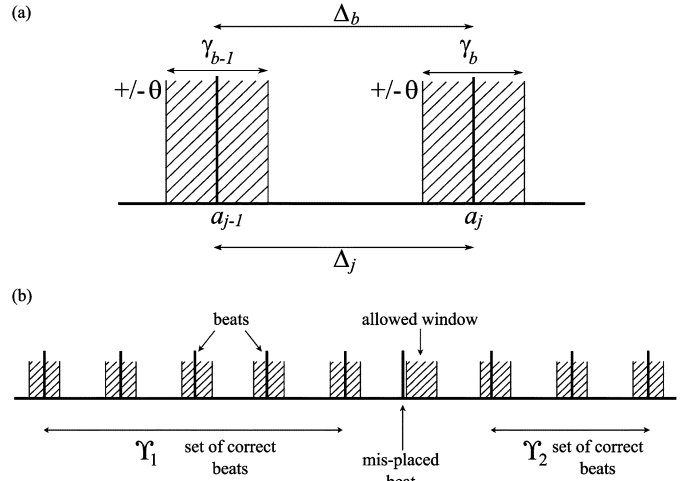


Fig. 6. (a) Continuity-based beat tracking evaluation metric. Acceptance of beat,  $\gamma_b$ , depends on its proximity to annotation  $a_j$  and the localization of the previous beat.  $\theta$  defines the acceptance window around each annotation. (b) Overview of evaluation metric over several beats.

The accuracy of given beat  $\gamma_b$  is defined by its proximity to annotation  $a_j$  and that of the preceding beat  $\gamma_{b-1}$  to annotation  $a_{j-1}$ . The dependence on the previous beat enforces the continuity constraint. We also require the current inter-beat-interval  $\Delta_b = \gamma_b - \gamma_{b-1}$  is consistent with the current inter-annotation-interval  $\Delta_j = a_j - a_{j-1}$ , creating the following three conditions:

- 1)  $a_j - \theta\Delta_j < \gamma_b < a_j + \theta\Delta_j$ ;
- 2)  $a_{j-1} - \theta\Delta_{j-1} < \gamma_{b-1} < a_{j-1} + \theta\Delta_{j-1}$ ;
- 3)  $(1 - \theta)\Delta_j < \Delta_b < (1 + \theta)\Delta_j$ ;

where in each condition,  $\theta$  defines the width of the acceptance window around each annotation. Klapuri *et al.* [13] evaluate their beat tracking using similar continuity-based criteria, and set  $\theta = 0.175$ , which equates to acceptance window of  $\pm 17.5\%$  around each annotation. The conditions for beat accuracy are shown Fig. 6(a).

For a sequence of beats  $\{\gamma\}$ , we can find the number of correct beats in each continuously correct segment  $\Upsilon_k$  [see Fig. 6(b)] that meet the criteria described previously

$$\hat{\Upsilon}_k = \max(\Upsilon_k) \quad k = 1, \dots, K \quad (41)$$

where  $K$  is the number of segments found. We extract  $\hat{\Upsilon}_k$  as the largest number of correct beats and therefore the longest continuously correct segment. To give a figure of overall accuracy for a given file, we find the ratio of  $\hat{\Upsilon}_k$  to the number of annotated beats  $\mathfrak{A}$  in that file

$$\text{accuracy}_{\text{cont}} = \frac{\hat{\Upsilon}_k}{\mathfrak{A}} \times 100\%. \quad (42)$$

The principle weakness of this approach, illustrated in Fig. 6(b), is that a single misplaced beat can cause the accuracy to drop significantly. In the worst case, where the erroneous beat is in the middle of the file, the accuracy can fall from 100% to 50%. To allow for continuity breaks, we also find the ratio of the total number of correct beats to the number of annotations

$$\text{accuracy}_{\text{total}} = \frac{\sum_{k=1}^K \Upsilon_k}{\mathfrak{A}} \times 100\%. \quad (43)$$



TABLE I  
COMPARISON OF BEAT TRACKING PERFORMANCE: CAUSAL ANALYSIS  
(C)—DP, KEA AND HUMAN DATA; NONCAUSAL (NC)—DP, KEA AND DIXON

Beat Tracker	CML cont. (%)	CML total (%)	AML cont. (%)	AML total (%)
DP (C)	<b>46.7</b>	<b>53.9</b>	<b>58.4</b>	<b>69.3</b>
KEA (C)	52.3	59.7	65.2	76.1
Human (C)	52.8	81.2	57.5	87.4
DP (NC)	<b>54.8</b>	<b>61.2</b>	<b>68.1</b>	<b>78.9</b>
KEA (NC)	55.7	62.4	70.0	80.0
Dixon (NC)	25.1	30.9	46.2	59.6

A further consideration is the choice of metrical level at which the tapping occurs, as humans often tap in time to music at different rates [25]. If we insist on the correct metrical level, then beats tapped at twice or half the rate of the annotations would be assigned an accuracy of 0% using (42) and (43), which may be considered to be too harsh. To allow for metrical ambiguity, we therefore also calculate a measure of continuously correct tracking, and the total number of correct beats for cases where tapping occurs at twice or half the annotated rate. To summarize, we present four measures of beat accuracy:

- correct metrical level, continuity required (CML cont);
- correct metrical level, continuity not required (CML total);
- allowed metrical levels, continuity required (AML cont);
- allowed metrical levels, continuity not required (AML total).

Results are shown in Table I first for causal analysis: our approach, labeled DP (C); the causal model of Klapuri *et al.* [13], labeled KEA (C); and the Human tapping performance; and then for noncausal analysis: DP (NC); the noncausal version of KEA (NC); and Dixon’s (noncausal) Beatroot method [1]. We find equivalent performance for our noncausal approach compared to that of KEA, both of which are more accurate than Dixon across each of the four measures of beat accuracy. We also observe the expected trend of poorer performance for the causal implementations of DP and KEA in comparison to the noncausal versions. Somewhat surprisingly, the Human tapping performance is weaker than both DP and KEA when continuity is required in the beat output.

Examining each column in Table I, we find weakest performance when continuity is required at the correct metrical level and strongest for the total number of correct beats at allowed metrical levels. This is only a partial ordering, however, as demonstrated by the results for the Human performance, where the accuracy for CML (total) is greater than that of AML (cont.). From this observation, we find the Human tapper to be most adversely affected by the continuity constraint. A second interesting feature of the Human performance relates to the relationship between results at CML and AML. In contrast to the computational approaches, which display a significant improvement when allowing for metrical ambiguity, the Human tapper only shows a moderate improvement, suggesting a more accurate inference of the annotated metrical level.

We can further examine the beat tracking accuracy as a function of the allowance window  $\theta$ . Results shown in Fig. 7 compare the noncausal approaches to the Human data under each of the four evaluation measures. As  $\theta > 20\%$ , the Human performance exceeds the algorithmic approaches under all condi-

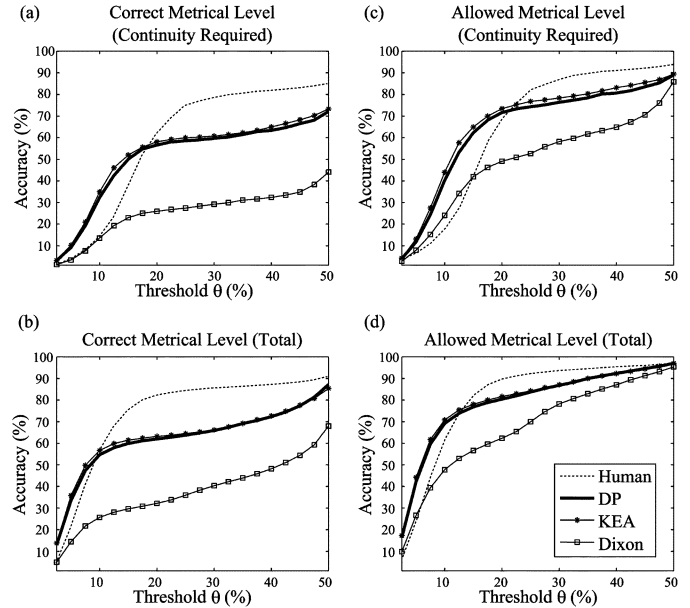


Fig. 7. Comparison of Davies (DP), Klapuri (KEA), Dixon, and Human beat tracking performance as a function of the acceptance window  $\theta$ . (a) Correct metrical level with continuity required. (b) Allowed metrical levels with continuity required. (c) Correct metrical level without continuity. (d) Allowed metrical levels without continuity requirement.

TABLE II  
BREAKDOWN OF PERFORMANCE ACROSS MUSICAL GENRE: CAUSAL ANALYSIS (C)—DP, KEA AND HUMAN DATA; NONCAUSAL ANALYSIS (NC)—DP, KEA, AND DIXON. RESULTS ARE MEASURED AT THE CORRECT METRICAL LEVEL WITH CONTINUITY REQUIRED

Beat Tracker	Dance (%)	Rock (%)	Jazz (%)	Folk (%)	Classical (%)	Choral (%)
DP (C)	<b>79.1</b>	<b>68.1</b>	<b>36.1</b>	<b>24.9</b>	<b>17.5</b>	<b>2.8</b>
KEA (C)	73.8	72.4	37.2	47.6	37.0	4.0
Human (C)	67.6	63.3	44.3	55.6	36.1	29.0
DP (NC)	<b>76.8</b>	<b>77.9</b>	<b>51.6</b>	<b>34.5</b>	<b>26.9</b>	<b>7.2</b>
KEA (NC)	76.3	77.9	39.1	52.2	40.2	4.0
Dixon (NC)	31.0	34.0	23.5	12.9	10.2	5.4

tions. In Fig. 7, it is clear exactly how close the performance of our proposed model DP is to KEA. The accuracy of the two approaches does not differ by more than 3% for all  $\theta$ . There are genre-related differences shown in Table II, where our noncausal approach is better able to infer the beats for Jazz examples, with KEA more accurate for Folk and Classical music. For Choral music, we find all algorithms having very low performance, with accuracy  $< 10\%$ . Only the Human tapper demonstrates any ability to find the beats.

### B. Beat Localization

An alternate way to visualize the performance of a beat tracker is to measure the *normalized beat error*—to show how close the beats are to the annotations. We find the normalized beat error,  $\epsilon_{n,j}$  for file  $n$ , as difference between each annotated beat  $a_j$  and the nearest generated beat  $\gamma_b$ , from all beats  $\gamma$ , divided by the inter-annotation-interval  $\Delta_j$

$$\epsilon_{n,j} = \min_{\gamma_b} \frac{|\gamma - a_j|}{\Delta_j}. \quad (44)$$

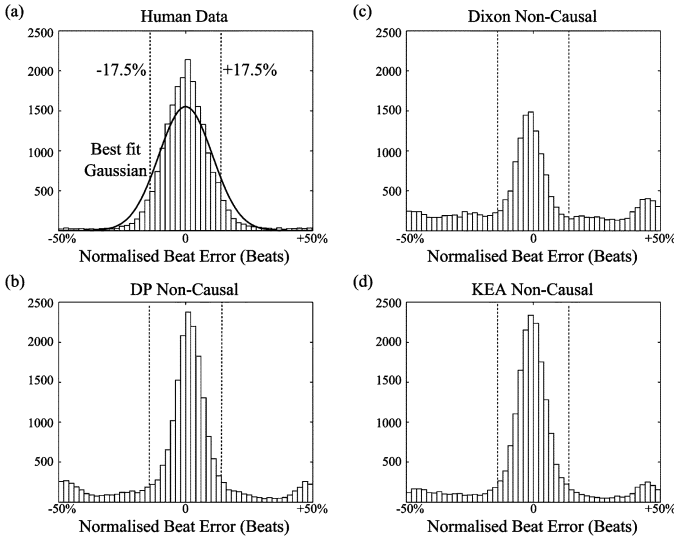


Fig. 8. Histogram of normalized beat error for (a) Human, (b) DP, (c) KEA, and (d) Dixon beat tracking algorithms. Vertical dotted lines indicate the acceptance window  $\theta$ .

At most, any one beat can be 50% ahead or behind the annotated value. This represents tapping in anti-phase to the annotated beats. By finding the error for all annotations in the database, approximately 20 000 beats, we can create histograms to show the distribution of beat error. These are shown in Fig. 8 for each of the noncausal approaches and the Human tapping data.

Each of the three algorithmic approaches [Fig. 8(b)–(d)] exhibit a strong peak close to an error of 0, indicating beats close to the annotations, and then two further peaks each around  $\pm 50\%$  of a beat away from the annotations, corresponding to tapping on the *off-beat*. We see an almost identical shape to the distributions of our approach and KEA. For Dixon’s algorithm, the distribution is flatter indicating a greater proportion of erroneous beats, which are neither close to 0, nor at  $\pm 50\%$ . The histogram of Human data has a different shape. It is close to Gaussian, with very few beats in anti-phase at  $\pm 50\%$  beats. This suggests greater consistency in finding the on-beat. However, the main peak of the distribution is wider than in the computational approaches indicating less precise beat localization.

### C. Computation Time

Having shown similar performance between our proposed model and the Klapuri *et al.* [13] approach, we now examine the computational complexity of these two algorithms. We have implemented our approach to mirror that of the KEA system, as generously supplied to us by Klapuri. Both systems are implemented in Matlab, but use precompiled functions to implement the generation of midlevel representation from the input audio signal. Due to the precompiled components, we are unable to quantify complexity in terms of mathematical operations (multiplications and additions) and therefore calculate the computation time ratio as the compute time (in seconds) for a file  $n$ , normalized by its length (in seconds), for all files  $N$  in the test database

$$\text{computation time ratio} = \frac{1}{N} \sum_{n=1}^N \frac{\text{compute time } (n)}{\text{length } (n)}. \quad (45)$$




Fig. 9. Proportion of normalized computation time taken in beat tracking compared to the generation of the midlevel representation for the KEA and DP algorithms.

Tests were run on a 3-GHz Linux machine running Matlab version 7.0. Results are given in Fig. 9 showing the time taken to generate the midlevel representation input and the remaining time to extract the beats from this representation. For our model, the computational complexity is significantly lower than for the KEA system. We also note that for our approach, the time taken in beat tracking is much less than the time required to generate the onset detection function midlevel representation. The converse is true for the KEA system, which spends a much more significant period of time recovering the beats.

## IV. DISCUSSION

We have shown our noncausal beat tracking algorithm to be of equivalent accuracy to the KEA system [13] when evaluated objectively using a continuity-based metric. Further, we have illustrated that this level of accuracy can be obtained at a considerable reduction in computational cost. By comparing the architecture of the two approaches, we can identify three factors in which the KEA system is more complex than our beat tracker: 1) Klapuri *et al.* [13] use a multiband approach to feature extraction, performing beat analysis in each subband; 2) for each analysis frame, the beat period and alignment are extracted simultaneously through parallel resonator filterbanks; and 3) KEA incorporate higher level knowledge using probabilistic analysis of multiple metrical levels. By making a single pass across one midlevel representation, separating the extraction of beat period and phase into two processes (therefore reducing a two-dimensional search into two, one-dimensional searches), and embedding high-level knowledge directly into the extraction of the beats, we are able to present a more efficient approach to beat tracking which is equally robust. Although parameter selection inevitably contributes to overall performance, those used in our beat tracker were based on sensible ranges for Western music. We have not attempted to find an optimal parameter set even though this could lead to our approach being more accurate than that of KEA, as we did not wish to over-train our beat tracker.

A notable weakness of our model in comparison to the KEA system is the extent to which accuracy is reduced for causal analysis. Tracking beats causally leads to a 10% reduction in accuracy compared to noncausal results, which for KEA is only a 3% reduction. It appears that our static prediction of future beats at precise beat period intervals (20) is not as effective as the higher level knowledge used within the KEA system. The simultaneous analysis of multiple metrical levels increases the reliability of beat predictions, which are directed towards the start

of the next measure. In our approach, we currently use no knowledge of the phase of higher metrical levels, extracting only the measure level periodicity. This is most telling in cases which exhibit expressive timing, where inaccurate beat predictions break the continuity requirement, causing a reduction in overall accuracy.

We see from the results in Table I that our causal approach is most competitive when continuity is not required and the metrical level is not specified (the AML total column). The Human performance is even more adversely affected by the requirement of a continuously correct output. We must question the importance of strict continuity when evaluating a beat tracking system. A single misplaced beat, in the context of a sequence of otherwise accurate beats, is a far less disturbing error than the accuracy calculated with the continuity requirement would suggest.

We have observed that the algorithmic beat trackers perform most accurately in cases where the beat structure is typically very rigid, for example in the Rock and Dance genres, where the human tapping often fails to give a continuously correct beat output especially when the allowance threshold  $\theta$  is less than 20% of a beat. Conversely, when  $\theta$  is increased, or the continuity requirement is relaxed, we see the accuracy of the human performance exceed all of the algorithmic approaches.

Three areas in which the human beat tracker can outperform the computational approaches are as follows: 1) in selecting the correct metrical level for tapping, as demonstrated by the small difference between results for CML and AML in Table I; 2) in selecting the correct phase, tapping on the on-beat rather than the off-beat, (see Fig. 8); and 3) tapping in time to music without any clear rhythmic structure, shown by the Choral examples in the test database (see Table II). These remain open challenges for automatic beat tracking, which may be solved by providing algorithms with higher level knowledge, including harmonic data, instrumentation, and genre. We plan to investigate these factors as part of our future work into rhythmic analysis of audio signals, in addition to the development of a real-time implementation of our approach toward an automatic musical accompaniment system.

## V. CONCLUSION

We have presented an efficient audio based beat tracking algorithm able to produce equivalent results to the current state of the art when evaluated using continuity-based criteria. The particular advantages of our model lie in its simplicity and computational efficiency, where our algorithm performs competitively at a significant reduction of the computational cost of the best performing published algorithm. In our future work, we intend to further investigate methods for evaluating beat tracking algorithms and extend our proposed model to operate in real-time for use within an interactive musical system.

## APPENDIX ONSET DETECTION FUNCTION

We use the complex spectral difference onset detection function as described in [17]. It is derived from analysis of overlapping short-term audio frames. To maintain a fixed time-resolution for the detection function that is independent of the sampling frequency of the input, we make the hop size  $h$  [in (46)]

dependent on the sampling frequency:  $h = f_s \times t_{\text{DF}}$ , where we set  $t_{\text{DF}}$  to represent the desired resolution (in seconds) of each detection function sample. We then use frames of length  $N = 2h$  creating a 50% overlap of the analysis at each iteration across the length of the input audio signal. For audio sampled at  $f_s = 44.1$  kHz, this creates a frame size of  $N = 1024$  with a hop size of  $h = 512$  audio samples.

We take the windowed short-term Fourier transform (STFT) of an input audio signal  $s(n)$  using an  $N$  length Hanning window  $w(m)$  and find the  $k$ th bin of the short term spectral frame  $S_k(m)$  as

$$S_k(m) = \sum_{n=-\infty}^{\infty} s(n)w(mh - n)e^{-j2\pi nk/N} \quad (46)$$

with a magnitude

$$R_k(m) = |S_k(m)| \quad (47)$$

where  $S_k(m) \in \mathbb{C}$  and  $R_k(m) \in \mathbb{R}$ .

By applying the assumption that during steady-state regions of the signal, i.e., not at a note onset, the magnitude spectrum should remain approximately constant, we can make a prediction of the magnitude spectrum  $\hat{R}_k$  at frame  $m$ , given the magnitude of the previous frame

$$\hat{R}_k(m) = R_k(m - 1). \quad (48)$$

In addition, we apply an assumption about the properties of the phase spectrum, that during steady state regions of the signal, the phase velocity at the  $k$ th bin should ideally be constant

$$\tilde{\phi}_k(m) - \tilde{\phi}_k(m - 1) \approx \tilde{\phi}_k(m - 1) - \tilde{\phi}_k(m - 2). \quad (49)$$

We adopt a short-hand notation for the left-hand side of (49)

$$\Delta\tilde{\phi}_k(m) = \tilde{\phi}_k(m) - \tilde{\phi}_k(m - 1). \quad (50)$$

By substituting (50) into (49) and rearranging terms, we can make a prediction about the phase of the  $k$ th bin for frame  $\tilde{\phi}_k(m)$  given the observations of the previous two frames

$$\hat{\phi}_k = \text{princarg}[\tilde{\phi}_k(m - 1) + \Delta\tilde{\phi}_k(m - 1)] \quad (51)$$

where  $\text{princarg}$  unwraps the phase value, mapping it into the range  $[-\pi, \pi]$ .

The predictions of the magnitude spectrum  $\hat{R}_k(m)$  and the phase spectrum  $\hat{\phi}_k(m)$  can be represented in polar form to give a spectral prediction  $\hat{S}_k(m)$  in the complex domain

$$\hat{S}_k(m) = \hat{R}_k(m)e^{j\hat{\phi}_k(m)} \quad (52)$$

which we compare to the observed complex spectrum

$$S_k(m) = R_k(m)e^{j\phi_k(m)}. \quad (53)$$

To derive the complex spectral difference detection function  $\Gamma(m)$  at frame  $m$ , we find the sum of the Euclidean distance between the predicted and observed spectra for all  $k$  bins

$$\Gamma(m) = \sum_{k=1}^K |S_k(m) - \hat{S}_k(m)|^2. \quad (54)$$

## ACKNOWLEDGMENT

The authors would like to thank A. Klapuri for making his source code for the KEA algorithm available, S. Hainsworth for providing the ground truth beat annotations, and N. Chétry for technical assistance.

## REFERENCES

- [1] S. Dixon, "Automatic extraction of tempo and beat from expressive performances," *J. New Music Res.*, vol. 30, pp. 39–58, 2001.
- [2] J. P. Bello, L. Daudet, S. Abdallah, C. Duxbury, M. Davies, and M. B. Sandler, "A tutorial on onset detection in music signals," *IEEE Trans. Speech Audio Process.*, vol. 13, no. 5, pt. 2, pp. 1035–1047, Sep. 2005.
- [3] F. Gouyon, A. Klapuri, S. Dixon, M. Alonso, G. Tzanetakis, C. Uhle, and P. Cano, "An experimental comparison of audio tempo induction algorithms," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 14, no. 5, pp. 1832–1844, Sep. 2006.
- [4] F. Gouyon and S. Dixon, "A review of automatic rhythm description systems," *Comput. Music J.*, vol. 29, no. 1, pp. 34–54, 2005.
- [5] J. P. Bello and J. Pickens, "A robust mid-level representation for harmonic content in music signals," in *Proc. 6th Int. Symp. Music Inf. Retrieval*, London, U.K., 2005, pp. 304–311.
- [6] M. A. Bartsch and G. H. Wakefield, "Audio thumbnailing of popular music using chroma-based representations," *IEEE Trans. Multimedia*, vol. 7, no. 1, pp. 96–104, Feb. 2005.
- [7] M. Levy, M. Sandler, and M. Casey, "Extraction of high level musical structure from audio data and its application to thumbnail generation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process (ICASSP)*, 2006, vol. V, pp. 13–16.
- [8] M. Goto, "An audio-based real-time beat tracking system for music with or without drum-sounds," *J. New Music Res.*, vol. 30, pp. 159–171, 2001.
- [9] S. Dixon, F. Gouyon, and G. Widmer, "Towards characterisation of music via rhythmic patterns," in *Proc. 5th Int. Conf. Music Inf. Retrieval*, Barcelona, Spain, 2004, pp. 509–517.
- [10] C. Uhle and C. Dittmar, "Drum pattern based genre classification of popular music," in *Proc. 25th Int. AES Conf. Semantic Audio*, London, U.K., 2004, pp. 189–195.
- [11] P. E. Allen and R. B. Dannenberg, "Tracking musical beats in real time," in *Proc. Int. Comput. Music Conf. (ICMC)*, Hong Kong, 1990, pp. 140–143.
- [12] E. D. Scheirer, "Tempo and beat analysis of acoustic musical signals," *J. Acoust. Soc. Amer.*, vol. 103, no. 1, pp. 588–601, 1998.
- [13] A. P. Klapuri, A. Eronen, and J. Astola, "Analysis of the meter of acoustic musical signals," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 14, no. 1, pp. 342–355, Jan. 2006.
- [14] D. F. Rosenthal, "Machine rhythm: Computer emulation of human rhythm perception," Ph.D. dissertation, School of Architecture and Planning, Mass. Inst. Technol., Cambridge, MA, 1992.
- [15] E. W. Large, "Beat tracking with a nonlinear oscillator," in *Working Notes of the IJCAI-95 Workshop on Issues in AI and Music*, 1995, pp. 24–31.
- [16] S. Hainsworth, "Techniques for the Automated Analysis of Musical Audio," Ph.D. dissertation, Dept. Eng., Cambridge Univ., Cambridge, U.K., 2004.
- [17] J. P. Bello, C. Duxbury, M. E. Davies, and M. B. Sandler, "On the use of phase and energy for musical onset detection in the complex domain," *IEEE Signal Process. Lett.*, vol. 11, no. 6, pp. 553–556, Jun. 2004.
- [18] J. C. Brown, "Determination of the meter of musical scores by auto-correlation," *J. Acoust. Soc. Amer.*, vol. 94, no. 4, pp. 1953–1957, Oct. 1993.
- [19] L. van Noorden and D. Moelants, "Resonance in the perception of musical pulse," *J. New Music Res.*, vol. 28, no. 1, pp. 43–66, 1999.
- [20] M. E. P. Davies and M. D. Plumbley, "Causal tempo tracking of audio," in *Proc. 5th Int. Symp. Music Inf. Retrieval*, Barcelona, Spain, 2004, pp. 164–169.
- [21] M. Minsky, "Music, mind and meaning," *Comput. Music J.*, vol. 5, no. 3, pp. 28–44, 1981.
- [22] F. Gouyon and P. Herrera, "Determination of the meter of musical audio signals: Seeking recurrences in beat segment descriptors," in *Proc. 114th Conv. Audio Eng. Soc.*, 2003, preprint 5811.
- [23] F. Gouyon, N. Wack, and S. Dixon, "An open source tool for semi-automatic rhythmic annotation," in *Proc. 7th Int. Conf. Digital Audio Effects (DAFx)*, Naples, Italy, 2004, pp. 193–196.
- [24] M. Goto and Y. Muraoka, "Issues in evaluating beat tracking systems," in *Working Notes of the IJCAI-97 Workshop on Issues in AI and Music—Evaluation and Assessment*, 1997, pp. 9–16.
- [25] D. Moelants and M. McKinney, "Tempo perception and musical content: What makes a piece fast, slow or temporally ambiguous?," in *Proc. 8th Int. Conf. Music Perception and Cognition*, Evanston, IL, 2004, pp. 558–562.



**Matthew Davies** (S'05) received the B.Eng. degree in computer systems with electronics from King's College London, London, U.K., in 2001. He is currently pursuing the Ph.D. degree in electronic engineering at Queen Mary University, London.

His current research interests include the rhythmic analysis of musical audio and interactive musical systems.



**Mark D. Plumbley** (S'88–M'90) received the Ph.D. degree in neural networks as a Research Student in the Engineering Department, Cambridge University, Cambridge, U.K., in 1991.

Following the Ph.D. degree, he joined King's College London, London, U.K., in 1991, and in 2002 moved to Queen Mary University, London, to help establish the new Centre for Digital Music. He is currently working on the analysis of musical audio, including automatic music transcription, beat tracking, audio source separation, independent component analysis (ICA), and sparse coding. He currently coordinates two U.K. research networks: the Digital Music Research Network (<http://www.dmrn.org>) and the ICA Research Network (<http://www.icarn.org>).

He is currently working on the analysis of musical audio, including automatic music transcription, beat tracking, audio source separation, independent component analysis (ICA), and sparse coding. He currently coordinates two U.K. research networks: the Digital Music Research Network (<http://www.dmrn.org>) and the ICA Research Network (<http://www.icarn.org>).