

Douglas Eck

Finding downbeats with a relaxation oscillator

Accepted: 14 June 2001 / Published online: 2 February 2002
© Springer-Verlag 2002

Abstract A relaxation oscillator model of neural spiking dynamics is applied to the task of finding downbeats in rhythmical patterns. The importance of downbeat discovery or ‘beat induction’ is discussed, and the relaxation oscillator model is compared to other oscillator models. In a set of computer simulations the model is tested on 35 rhythmical patterns. The model performs well, making good predictions in 34 of 35 cases. In an analysis we identify some shortcomings of the model and relate model behavior to dynamical properties of relaxation oscillators.

Introduction

The term *beats* refers to sounds that are perceived as being equally spaced in time. *Downbeats* are particularly salient beats that usually occur at a comfortable tapping rate. When you tap your feet to the radio you are finding downbeats, a skill called *beat induction*. Downbeats act as a unifying force, lending music the feeling of movement by allowing the listener to predict the onset of important musical events. The process of beat induction is influenced by many aspects of music including harmony, melody and rhythm (Cooper & Meyer, 1960; Lerdahl & Jackendoff, 1983). Because interactions are not always simple, it can be difficult to predict the locations of downbeats. For example, in rock and roll music, even though chord changes (harmonic components) usually occur on the first note of a musical bar, downbeats are often aligned with the second note due to syncopated drumming style. Furthermore, although beats are *perceived* as being equally-spaced in time, perfect timing is rarely if ever found on the radio. Both

motor noise and deviations due to expressive timing are found in performed music.

The task of beat induction can be simplified by considering only patterns of equal-amplitude beeps or clicks. This removes the influence of melody and harmony (and in addition amplitude and timbre variations). However, even in this simpler domain, two important influences compete to decide downbeat location. First, the *grouping* of events in the pattern is important. For example, when three or more events are presented in rapid succession, the first and last events are more perceptually salient than those in the middle (Povel & Okkerman, 1981). Second, the *meter* of the pattern is important. Meter is the sense of strong and weak beats that arises from interactions among hierarchical levels in a pattern having nested periodic components. Such a hierarchy is implied in Western music notation, where different levels are indicated by kinds of notes (whole notes, half notes, quarter notes, etc.) and where bars establish measures of an equal number of beats. For an overview see Handel (1989).

Although beat induction is relatively simple for people to perform – most of us can tap our feet to a child’s playground song – it is deceptively difficult to model computationally. We address this challenge by presenting a dynamical model of beat induction. Our model uses a nonlinear oscillator characterized by alternating slow and fast movement. This so-called *relaxation oscillator* is used to model a range of biological oscillatory processes, including neural spiking (Hodgkin & Huxley, 1952) and heartbeat pacemaking (Van der Pol & Van der Mark, 1928). Here we motivate the use of relaxation oscillators in the domain of beat induction and compare our approach to other oscillator approaches. After describing our model in some detail we show, through a series of simulations, that a relaxation oscillator can perform beat induction in a way that (to an extent at least) mirrors human performance. In analyzing the simulation results we identify some shortcomings of the model and relate model behavior to some relevant dynamical properties of relaxation oscillators.

D. Eck
Istituto Dalle Molle di Studi sull’Intelligenza Artificiale (IDSIA),
Galleria 2, 6928 Manno-Lugano, Switzerland
E-mail: doug@idsia.ch
URL: www.idsia.ch/~doug

Oscillators and beat induction

Oscillators have been applied to beat induction and similar rhythm cognition tasks for at least 15 years. Most have been *limit cycle* oscillators, meaning that when perturbed they quickly return to the same path or cycle in phase space. This is advantageous because it allows an oscillator to be perturbed by, for example, a note onset in a rhythm without loss of stability: even after withstanding a very strong perturbation the oscillator is guaranteed to return to the limit cycle. If the limit cycle is nonlinear, perturbations at some points in the limit cycle will have more effect on the oscillator than others. Using various strategies, this behavior can be exploited so that perturbations to an oscillator cause it to align itself with some periodic component. That is, nonlinear response to perturbations can lead to an oscillator naturally *beating along* with driving signals having compatible frequencies.

Dannenberg (1984) used such an oscillator to find downbeats in patterns and to track acceleration and deceleration by modifying oscillator period based on changes in pattern rate. Torras (1985) used firing threshold adaptation in limit oscillators similar to those used here; the task was to find temporal regularities in simple rhythms. Miller, Scarborough, and Jones (1992) used a coupled one-dimensional network of oscillators (BEAT-NET) to resonate with rhythmical patterns. There are also examples of models where nonlinear oscillation is a component of a cognitive model of rhythm, but is not a primary part of the system. For example, Todd, O'Boyle, and Lee (1999) incorporate an oscillator model of musculoskeletal movement in a system that synchronizes body movements with temporal regularities in an input signal.

Similar oscillator models by McAuley (McAuley, 1994) and Large (Large & Jones, 1999; Large & Kolen, 1994) are successful at finding downbeats in patterns even when non-stationary noise (e.g., acceleration) is present in the patterns. McAuley used the term *adaptive oscillator* to describe a limit cycle oscillator that entrains both its phase and its period to recurring events in a temporal signal. The McAuley oscillator entrains to a rhythmical pulse train by discretely resetting its phase to zero when a pulse is sufficiently strong. This phase resetting is governed by a cosine-shaped function that is centered around the zero phase of the oscillator. By tightening the shape of this function, the system has the ability to focus on a particular periodic component of the signal, ignoring all others. The McAuley adaptive oscillator also attempts to match its period to periodic components in the signal. This is achieved by a function that slightly slows the oscillator when phase resetting consistently occurs early, and slightly accelerates the oscillator when phase resetting consistently occurs late.

Large and Kolen (1994) proposed a nonlinear limit cycle oscillator that entrains its phase to a rhythmical input by means of gradient descent. A smooth function

that crosses zero at phase zero is used to establish a phase-based attractor. That function is minimized with respect to the difference between oscillator phase and the occurrence of input events, resulting in an oscillator that continually aligns its zero phase with that of events in the input. The Large oscillator has a second variable that modifies the slope and width of the gradient descent function such that the oscillator can sharpen its receptive field, allowing it to lock onto specific periodic components in the signal. Large and Kolen (1994) show that such an oscillator can form the basis of small connectionist networks that find salient events at multiple levels of the metrical hierarchy.

Research goals

It is clear, even with this cursory sampling of models, that many approaches to oscillator beat induction have already been attempted. Why try another one? Our goal is not to contribute to a "wild west" situation of divergent models. Instead, we wish to start with a model having some interesting dynamical properties and understand the extent to which it can find downbeats in patterns. In other words, we are not so much building another oscillator beat tracker, as borrowing an oscillator from neurobiology and applying it to the task.

Our motivation for choosing a relaxation oscillator stems from research showing that relaxation oscillators have better synchronization properties than similar single-time-scale limit cycle oscillators. Namely, when large numbers of relaxation oscillators are coupled together, they readily synchronize their oscillation. Non-relaxation oscillators are much less able to achieve this group synchrony. Somers and Kopell (1993, 1995) describe this behavior in detail and offer a theory called Fast Threshold Modulation (FTM) that attributes this behavior to the modulation of firing thresholds in response to input voltage. The authors show that relaxation oscillators exhibit FTM but that non-relaxation oscillators in general do not, and they use this point to argue for the superiority of relaxation oscillators in tasks requiring robust synchronization. The details of their argument are not important for our purposes, and readers are referred to the original sources for an explanation. What is important is the observation that relaxation oscillators are very good synchronizers.

This observation is important for two reasons. First, synchrony is the means by which an oscillator finds downbeats in a pattern. Thus, it is not unreasonable to think that an oscillator with good general synchronization dynamics may excel at this task. Second, although only a single oscillator is examined in this study, one of our goals is to use large networks of coupled oscillators for storing and retrieving entire rhythmical patterns (Eck, 1999, 2001). In this case, group synchrony among many coupled oscillators is vital; without it, a network would be unstable. In previous attempts at performing

network learning (Gasser, Eck, & Port, 1999), we noted that a single-variable limit cycle oscillator similar to McAuley's and Large's oscillator worked well in small ensembles, but that groups larger than four or five oscillators did not work because the networks would not synchronize and so became unstable. Thus, group synchronization seems to be an important consideration when scaling up to the many oscillators required for network learning.

One component missing from this study is a comparison of several models on the same set of patterns. Although this was our original goal, in the end we were unable to achieve success with some other models. We do not believe this shows our model to be superior. Rather, we suspect that even by systematically trying different parameter settings and input encodings, we did not arrive at the right ones. As our model is very sensitive to certain parameter values (input amplitude for example), we do not consider this a criticism of the other models. Note that we made every attempt to clearly define how we encoded the input pattern set and how we measured performance of our model. This makes it possible for other researchers to run comparison simulations if they wish.

A relaxation oscillator model of beat induction

The Fitzhugh-Nagumo relaxation oscillator (Fitzhugh, 1961; Nagumo, Arimoto, & Yoshizawa, 1962) is a two-variable model of neural action potential. It is a simplification of a more complicated model by Hodgkin and Huxley (1952). Under a wide range of parameter settings the Fitzhugh-Nagumo oscillator exhibits the spiking dynamics of a real neuron: it gradually accrues voltage until it reaches a threshold; upon reaching that threshold it fires and quickly releases the energy. With constant and sufficient driving energy, this results in stable limit cycle oscillation.

This behavior is achieved by coupling two equations together, with one modeling slow uptake using a cubic function and the other modeling fast release using a linear function. The nullclines for these functions are seen in Fig. 1 in the phase portrait on the left. These are computed by setting $dv/dt=0$ (generating the voltage v nullcline) and $dw/dt=0$ (generating the voltage recovery w nullcline) in the equation for the Fitzhugh-Nagumo oscillator:

$$\begin{aligned}\frac{dv}{dt} &= -v(v - \theta)(v - 1) - \omega + \Omega \\ \frac{dw}{dt} &= \epsilon(v - \gamma\omega)\end{aligned}\quad (1)$$

The variable ϵ is used to control the frequency of oscillation and shape of waveform: a high value yields high-frequency oscillation with a sinus-like waveform, while a low value yields low-frequency relaxation oscillation. The variable Ω represents driving energy. In the simple case Ω is the small amount of constant external voltage

necessary for sustained oscillation. For our purposes we add to this small amount of voltage a rhythmical input pattern encoded as voltage spikes. The two remaining parameters are fixed for all simulations. They include a firing threshold parameter θ set to 0.2, and a shunting parameter γ set to 1.2.

So that we can easily test model behavior, we modified the Fitzhugh-Nagumo oscillator so that it generates a discrete output event with every cycle. These output events were only used to measure synchrony; they played no role in simulation. As there is no predefined point in the limit cycle to mark output, (i.e., no phase zero) we had to choose one. The value where voltage is equal to 0.4 and rising was used because it lies between the approximate voltage minimum of -0.2 and voltage maximum of 1.0 of the oscillator limit cycle. It can be seen on the right in Fig. 1 that this value did a good job in marking the onset of a spike, although near values would have worked as well.

To sum, we started with a relaxation oscillator model of neural action potential having good synchronization properties. We modified it so that we could drive the oscillator with rhythmical patterns. Finally, we marked a point in the limit cycle corresponding to neural firing so that we could easily compute the alignment of an oscillator with events in the input.

Simulations

To test the model we used a set of rhythms from Experiment 1 of Povel and Essens (1985). These rhythms are generated by permuting the interval sequence **1 1 1 1 1 2 2 3** and terminating it by the interval **4**. These length-16 patterns all contain nine notes and seven rests, and are cycled for the oscillator. In Table 1 the patterns are seen ordered by row. The numbers in the first column correspond to difficulty rankings (from 1 to 7) based upon a non-oscillator model of downbeat induction by Povel and Essens from the same paper.

Method

Input patterns

The patterns in Table 1 were transformed for the oscillator by generating a vector of zeros and inserting a 1 whenever a note onset was to occur. So that the oscillator would run in low-frequency relaxation mode, the base interval between notes was 125 time steps. Thus, an x in Table 1 was transformed into a 1 followed by 124 zeros. A $.$ was transformed into 125 zeros. The amplitude of the signal was then modulated so that the input signal started relatively low and increased evenly throughout the duration of the simulation. This strategy of steadily increasing input amplitude kept the oscillator from always preferring early notes in the sequence. Early-note preference may seem like a weakness in our model, but in fact listeners exhibit the same preference (Garner & Gottwald, 1968).

Two sets of simulations were performed based on input signal strength. For the first set (low signal strength) the input pulse voltage started at 0.0625 and increased linearly to a maximum 0.08 (mean = 0.07125). For the second set (high signal strength) the input pulse voltage started at 0.0625 and increased to 0.09 (mean = 0.0775).

Flow diagram and time series




Fig. 1 The phase portrait (*left*) shows the limit cycle of the oscillator as it follows the voltage nullcline. The rate of movement is fast along the *horizontal* (voltage) axis and slow along the *vertical* (recovery) axis. This slow-fast relaxation oscillation is also visible in the time series (*right*) showing the oscillator voltage waveform. Oscillator output events are marked with *triangles* in both diagrams

To look at sensitivity to noise in the input signals, two sets of noisy signals were used in addition to the noise-free ones. Noise was injected by adjusting the inter-onset interval (IOI) between pulses. Noise was not injected into the amplitude of the input signal, although this is an interesting avenue for further simulations. For the low noise set, all IOIs were chosen uniform randomly from the range -5% to $+5\%$ of the base IOI 125 (118.75 to 131.25). For the high noise set, the range was -10% to $+10\%$ (112.5 to 137.5). Non-stationary transformations of the input such as acceleration and deceleration were not tried, although these are also of future interest.

Two variations in signal strength and three variations in noise yielded six separate simulations sets. All simulations were run using the ordinary differential equation (ODE) solver suite in Matlab 5.3. Several ODE solvers were tested to ensure that computational errors were not an issue.

Simulation details

Although the oscillator model is a deterministic one, it makes different predictions based on its initial conditions. For example, if the initial location of an oscillator in phase space (i.e., initial (v, w) setting) causes alignment with some recurring event in the signal, the oscillator will phase lock with that event at the expense of finding other recurring events. For this reason, the best way to understand the model is to simulate many oscillators having different initial conditions.

For these experiments 20 oscillators were used. To vary performance, starting locations were evenly distributed high on the left-hand leg of the cubic voltage nullcline (found in the upper left-hand quadrant of the phase diagram, left in Fig. 1). Because

movement down this nullcline is slow, the oscillators reached the limit cycle and began oscillating at well-spaced intervals. All 20 oscillators were computed in parallel, but they were not coupled and had absolutely no effect on one another.

The voltage-to-recovery coupling value (ϵ in Equation 1) for all oscillators was set such that the period of oscillation was four times slower than the base IOI of the pattern (125). This yielded a group of period 500 oscillators ($\epsilon = 0.0015$). By setting the period of the oscillators in this way, the task was simplified because the oscillators did not need to use period adaptation to find the eigenfrequencies in the pattern. However, this choice did not trivialize the problem because the oscillator must still synchronize with the input. That is, the oscillator must be stable enough to respond to some notes in the input (those indicating the downbeat), while ignoring all others. Furthermore, we believe the issue of period adaptation may perhaps be solved by a group of oscillators having a range of intrinsic periods rather than by a single oscillator that continually adapts its own period.

Measurement

Downbeat induction was measured using a binning method that matched oscillator outputs to locations in the input pattern. After allowing the system to run for eight pattern repetitions, oscillator output for two following pattern repetitions were assigned a phase relative to the input pattern. Since the oscillator was running four times slower than the base input pattern IOI, there were four possible phases, numbered 1 through 4 from the beginning of a pattern (Table 1) in later analyses.

If for a particular oscillator the bins did not match for both pattern repetitions, that oscillator made no downbeat prediction. Thus, successful oscillators are those which always fired at the same phase and so revealed stable periodic oscillation. Through a large number of simulations, this method was shown to be an effective measure of beat induction. Namely, using more than eight settling repetitions did not yield a significant increase in success, while more than two binning repetitions did not yield a significant increase in failure.

Table 1 Thirty-five patterns from Povel and Essens (1985) ordered by row. In these patterns, an “x” indicates a note and a “.” indicates a rest. The difficulty ratings in the *left-hand column* are from a non-oscillator model in the same paper and range from easy (1) to hard (7)

1	xxxxx..xx.x.x...	xxx.x.xxx..xx...	x.xxx.xxx..xx...	x.x.xxxxx..xx...	x..xx.x.xxxxx...
2	xxx.xxx.xx..x...	x.xxxx.xx..xx...	xx..xxxx.x.x...	xx..x.xxx.xxx...	x.xxx.xxxx..x...
3	xxx.xx..xx.xx...	xx.xxxx.x..xx...	xx.xx.xxxx..x...	xx..xx.xx.xxx...	x..xxx.xxx.xx...
4	xx.xxxx.xx..x...	xx.xxx.xxx..x...	xx.xxx..xx.xx...	xx..xx.xxxx.x...	xx..xx.xxx.xx...
5	xxxxx.xx.x..x...	xxxx.x..xxx.x...	xxx..xx.xxx.x...	x.xxx..x.xxxx...	x.x..xxxx.xxx...
6	xxxx.x.x..xxx...	xx.xxx.x..xxx...	xx.x..xxx.xxx...	x.xxxx.x..xxx...	x..xxxxx.xx.x...
7	xxxx.xxx..x.x...	xxxx..xx.xx.x...	xx.xxxx..xx.x...	xx.x.xxxxx.x...	x.x..xxx.xxxx...

Results

Due to space constraints, the simulation results are summarized here. The full set of results are available at www.idsia.ch/~doug/publications.html. Audio examples are also found at the same site. Table 2 provides an overall summary of the performance data from the simulations.

Beat induction is a perceptual phenomenon that changes with musical experience and pattern presentation rate (Duke, 1989; Parncutt, 1994; McAuley & Semple, 1999). For this reason there is no single correct downbeat assignment for a given pattern, and performance is best measured by comparing model predictions to a large set of human experimental beat assignments. Unfortunately, to our knowledge, no such dataset exists. Povel and Essens (1985) performed several experiments using these patterns, but the experiments concern pattern learning and recall, not downbeat assignment. For this analysis we compared the oscillator downbeat inductions to the predictions made by the Povel and Essens' (1985) P&E model (see Eck (in press) for our treatment of the model). We chose this model because of its success at predicting errors in learning and recall for this pattern set.

With P&E downbeat predictions as our criterion for success, we report that the Fitzhugh-Nagumo oscillator makes correct downbeat predictions for 34 of 35 patterns, a success rate of 97.1%. Pattern 28 was the only total failure. In two of the five cases where there were multiple correct downbeat assignments made by the P&E model, the oscillator found both of them (patterns 23 and 33). In three cases (patterns 22, 25 and 29), the oscillator model predicted two different downbeats, while the P&E model predicted only one. Overall the performance of the model was very good. However, a caveat is in order: for some difficult patterns, a large percentage of oscillators failed to settle into periodic oscillation. We discarded those oscillators when making downbeat predictions, in some cases choosing a downbeat from a very small subset of the 20 oscillators. Failure rate is discussed in more detail below. Readers are referred to the full dataset on the website for more details.

Two examples of the performance of the model are shown in Fig. 2. The top graph shows the oscillators

tracking pattern 1. All 20 oscillators have settled on the same solution of Phase 1. The bottom graph shows pattern 23 where the oscillators have found two plausible downbeats at Phases 1 and 2.

Discussion

The effects of input signal noise

Injecting noise into signals caused general degradation in performance but caused neither catastrophic failure nor differential behavior. In our view, this shows that the oscillator is robust to noise. For example, even with 20% noise, the oscillator was able to find 29 of 35 downbeats (82.9%) with low-amplitude input signals (see Table 2). Recall that this noise was stationary. Non-stationary transformations of the input such as acceleration and deceleration were not investigated but would be good candidates for future research.

The effects of input signal strength

Although noise did not pose a problem for the model, variations in input strength did. In Table 2 it can be seen that success rate drops significantly for the higher amplitude, especially with noise. By closely analyzing beat assignments, it was discovered that the increase in input strength causes the oscillator to show an over-sensitivity to *early* events (i.e., the first event in a series of connected events, as well as isolated events). In Fig. 3, pattern 15 is used as an example to show this behavior. With low signal strength (top) all oscillators predict a Phase 1 downbeat, the correct prediction with respect to the P&E model. With high signal strength (bottom), however, 13 of 20 oscillators move to the incorrect Phase 4 solution, with only three oscillators settled on the Phase 1. Note that in the high strength Phase 4 solution the oscillator has aligned with three *early* events, while in the low strength Phase 1 solution, there is alignment with only one.

This sensitivity to early events can be explained in terms of relaxation oscillator dynamics: when signal strength is very strong, the first event in a series will cause the oscillator to fire. During recovery, the oscillator is unable to respond to subsequent (perhaps more rhythmically important) events.

Oscillator failure as predictor of pattern complexity

As was already mentioned, those oscillators which did not settle into periodic oscillation were treated as failures by the binning algorithm. The rate of failure was not the same for all patterns. In fact, it varied considerably. In Fig. 4, failure rate is used as a predictor of relative pattern complexity, with a high rate of failure indicating a more complex pattern to track. As can be

Table 2 Summary of simulation results. The values presented are number and percent of successful downbeat predictions. Success was measured by comparing oscillator predictions to predictions of the Povel and Essens' (1985) P&E model. The *columns* indicate signal strength (*low* and *high*). The *rows* indicate noise amount (none, 10%, 20%)

Simulation summary		
Noise	Low strength	High strength
None	34 (97%)	28 (80%)
10%	32 (91%)	7 (20%)
20%	29 (83%)	8 (23%)

Fig. 2 The time series of 20 oscillators tracking pattern 1 (*top*) and pattern 23 (*bottom*). Eight pattern presentations are shown. The input pattern (*inverted triangle*) and oscillator output (*triangle*) events can be seen above the oscillator voltage waveforms. For pattern 1, all oscillators settled on a single downbeat assignment. For pattern 23, two solutions were found

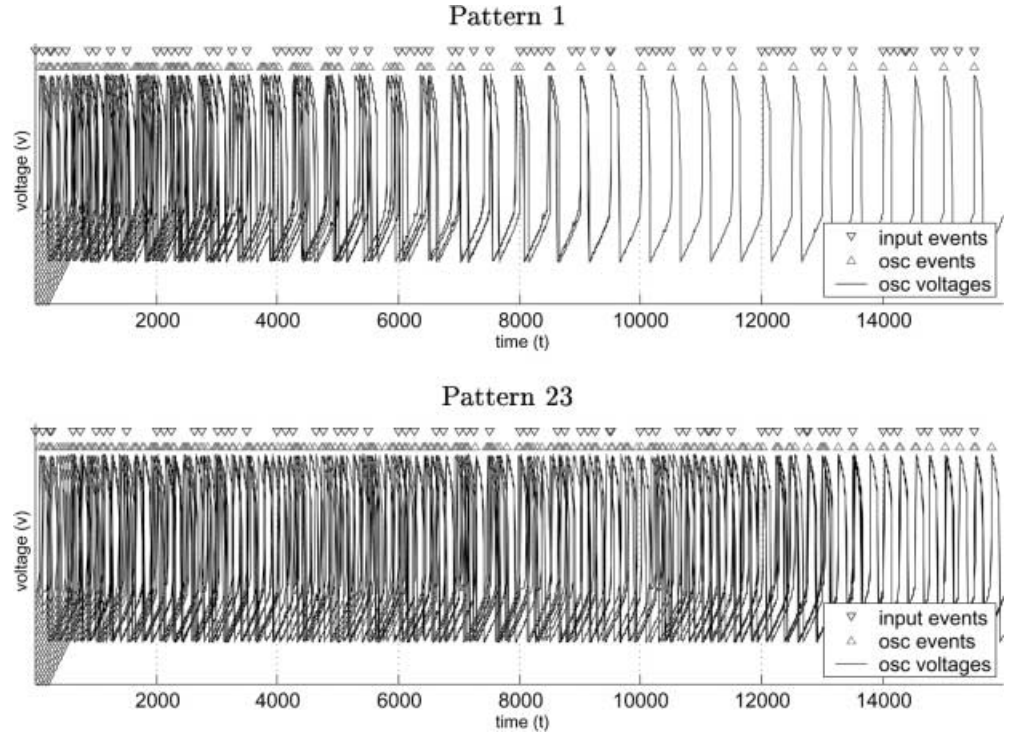
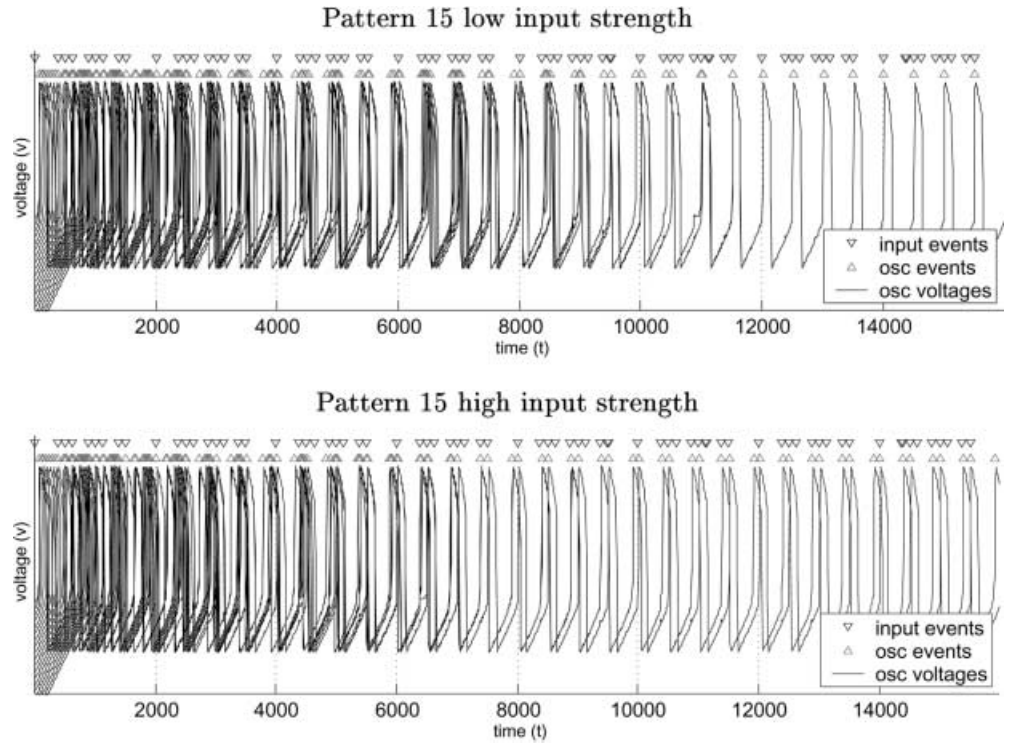


Fig. 3 On the *top* is the pattern 15 presentation with low input strength. Here, all oscillators align with the predicted Phase 1 downbeat. On the *bottom* is the pattern 15 presentation with high input strength. Here, some oscillators settle at Phase 4, a solution not predicted by the P&E model (Povel & Essens, 1985). The high input strength causes the oscillators to be overly sensitive to early events in the input



seen, the failure rates for the oscillator model jump considerably at Pattern 21. An analysis of pattern structure reveals that this is the point in the pattern set where rests become a necessary component in any downbeat assignment. This is the case because for patterns 1 through 20 at least one period-four downbeat

solution encounters only beats, while for patterns 21 through 35 no period-four downbeat solution encounters only beats. Recall that all patterns are terminated by an interval 4 (a note followed by three rests). Thus, for all patterns, three of four possible period-four downbeat solutions will encounter a rest. The transition at pattern




Fig. 4 The number of failed oscillators for each pattern is plotted, along with normalized complexity predictions from the P&E model. Two P&E weights (W) are shown. This weight adjusts the relative importance of rests in ranking pattern complexity. The performance of the oscillator model is more in accordance with P&E when weight W is unrealistically high ($W=10$). This indicates that the oscillator mechanism is over-sensitive to the presence of rests, a weakness in the model

21 occurs when a rest is introduced for the remaining possible solution.

Along with oscillator failure rate, Fig. 4 also shows two different sets of P&E complexity rankings. The set that gradually stair-steps up from left to right (diamonds) places moderate importance on the presence of rests. This ranking (P&E importance-of-rests parameter $W=4$) provided a good match to subject ability to memorize and recall the patterns (Povel & Essens, 1985). The set that drastically increases in complexity at pattern 21 (Xs) is generated using an unrealistically high ($W=10$) weighting of the importance of rests. Note that our oscillator performance more closely matches this complexity ranking. This is another way of showing that the oscillator model is overly sensitive to rests in a pattern.

This sensitivity can be explained in terms of relaxation oscillator dynamics. An incoming note (an input voltage spike) perturbs the oscillator, often causing it to fire. However, a rest has no perturbing energy and so is unable to force firing to occur. For this reason, an oscillator synchronized with a particular periodic component in a pattern may fire slightly late upon encountering a rest and fall slightly behind. If an oscillator falls far enough behind to fire in response to a *subsequent* beat, catastrophic failure ensues: the oscillator is now completely aligned with this new subsequent beat. This is the case because a relaxation oscillator does not *gradually* adjust its phase in response to input but instead fires completely. This is, in part, the same behavior that makes the oscillator good at synchronizing in groups. However, in this case it leads to synchronization behavior that is unfavorable for beat induction.

Future research

First, a good comparison of several oscillator models needs to be undertaken. Although we remarked that such a comparison was too difficult for this study, we believe that with collaboration between researchers such a survey is possible. Second, we are aware of no comprehensive set of listener downbeat assignments for the Povel and Essens patterns or similar patterns. The collection of such data would be of great value. Furthermore, we believe that if such a study is undertaken, additional patterns should be included that explicitly test the relative salience of *early* versus *late* beats (e.g., the first in a series of three notes versus the last in the series). Our model predicts that early events should be more salient, a prediction we currently have not tested.

Recent experiments show that the observed oversensitivity to rests is less severe in coupled *networks* of relaxation oscillators. These early results suggest that oscillator-to-oscillator coupling makes the network more stable and able to ‘keep the beat’ even in the presence of rests (see Eck, 2001 for these results). We believe more research is necessary in this direction.

Finally, the current study has been concerned only with synchronization. It ignores ‘continuation phase’ behavior; that is, the behavior of the beat induction mechanism when the input pattern is removed. In fact, it could be said that the oscillator is only ‘following’ because it never generates predictions in the absence of an input signal but rather only responds to it. We believe that the oscillator is doing a very special kind of *following* by synchronizing with downbeats, while ignoring less salient events in the signal. However, we agree that this question of relaxation oscillator behavior in continuation deserves attention.

Conclusions

This study, we believe, establishes a relaxation oscillator as a good candidate for beat induction. It is able to find downbeats in a number of patterns, many of them reasonably complicated. In this way we have shown that a relaxation oscillator is able to focus on a single periodic component in a complicated pattern (the downbeat), while ignoring all others. Many previous oscillator approaches have required special input filtering functions to achieve the same success. Although more research is clearly needed, we believe that the current study has successfully extended the findings on robust relaxation oscillator synchrony (Somers & Kopell, 1995) into the domain of rhythmical beat induction.

Acknowledgements Thanks to Nancy Kopell, David Somers and Art Winfree for insights on oscillator theory. Thanks to anonymous reviewers for many helpful suggestions. This research was supported by SNF project 21-49144.96.

References

- Cooper, G., & Meyer, L. B. (1960). *The rhythmic structure of music*. Chicago: University of Chicago Press.
- Dannenberg, R. (1984). An on-line algorithm for real-time accompaniment. In *Proceedings of the 1984 International Computer Music Conference*. Computer Music Association.
- Duke, R. A. (1989). Musicians' perception of beat in monotonic stimuli. *Journal of Research in Music Education*, 37, 61–71.
- Eck, D. (1999). Learning simple metrical preferences in a network of Fitzhugh-Nagumo oscillators. In *The Proceedings of the Twenty-First Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Eck, D. (2001). A network of relaxation oscillators that finds downbeats in rhythms. In *Proceedings of ICANN 2001, International Conference on Artificial Neural Networks*, Vienna, Austria, 2001. London: IEE.
- Eck, D. (in press). A positive-evidence model for rhythmical beat induction. *Journal of New Music Research*.
- Fitzhugh, R. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1, 455–466.
- Garner, W., & Gottwald, R. (1968). The perception and learning of temporal patterns. *Quarterly Journal of Experimental Psychology*, 20, 97–109.
- Gasser, W., Eck, D., & Port, R. (1999). Meter as mechanism: a neural network model that learns metrical patterns. *Connection Science*, 11, 187–216.
- Handel, S. (1989). *Listening: an introduction to the perception of auditory events*. Cambridge, MA: MIT Press.
- Hodgkin, A., & Huxley, A. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117, 500–544.
- Large, E. W., & Jones, M. R. (1999). The dynamics of attending: how people track time-varying events. *Psychological Review*, 106, 119–159.
- Large, E. W., & Kolen, J. F. (1994). Resonance and the perception of musical meter. *Connection Science*, 6, 177–208.
- Lerdahl, F., & Jackendoff, R. (1983). *A generative theory of tonal music*. Cambridge, MA: MIT Press.
- McAuley, J. D. (1994). Finding metrical structure in time. In M. C. Mozer, P. Smolensky, D. S. Touretsky, J. L. Elman & A. S. Weigend (Eds.), *Proceedings of the 1993 Connectionist Models Summer School* (pp. 219–227). Hillsdale, NJ: Erlbaum.
- McAuley, J. D., & Semple, P. (1999). The effect of tempo and musical experience on perceived beat. *Australian Journal of Psychology*, 51, 176–187.
- Miller, B. O., Scarborough, D. L., & Jones, J. A. (1992). On the perception of meter. In M. Balaban, K. Ebcioglu & O. Laske (Eds.), *Understanding music with AI: perspectives on music cognition* (pp. 429–447). Cambridge, MA: MIT Press.
- Nagumo, J., Arimoto, S., & Yoshizawa, S. (1962). An active pulse transmission line simulating nerve axon. *Proceeding IRE*, 50, 2061–2070.
- Parncutt, R. (1994). A perceptual model of pulse salience and metrical accent in musical rhythms. *Music Perception*, 11, 409–464.
- Povel, D., & Essens, P. (1985). Perception of temporal patterns. *Music Perception*, 2, 411–440.
- Povel, D., & Okkerman, H. (1981). Accents in equitone sequences. *Perception and Psychophysics*, 30, 565–572.
- Somers, D., & Kopell, N. (1993). Rapid synchronization through fast threshold modulation. *Biological Cybernetics*, 68, 393–407.
- Somers, D., & Kopell, N. (1995). Waves and synchrony in networks of oscillators of relaxation and non-relaxation type. *Physica D*, 89, 169–183.
- Todd, N., O'Boyle, D., & Lee, C. (1999). A sensory-motor theory of rhythm, time perception and beat induction. *Journal of New Music Research*, 28, 5–28.
- Torras, C. (1985). *Temporal-pattern learning in neural models*. Berlin: Springer.
- Van der Pol, B., & Van der Mark, J. (1928). The heartbeat considered as relaxation oscillation. *Phil. Mag. Suppl.*, 6, 763–775.