

Representing Rhythmic Patterns in a Network of Oscillators

Michael Gasser and Douglas Eck
Indiana University

Abstract

This paper describes an evolving computational model of the perception and production of simple rhythmic patterns. The model consists of a network of oscillators of different resting frequencies which couple with input patterns and with each other. Oscillators whose frequencies match periodicities in the input tend to become activated. Metrical structure is represented explicitly in the network in the form of clusters of oscillators whose frequencies and phase angles are constrained to maintain the harmonic relationships that characterize meter. Rests in rhythmic patterns are represented by explicit rest oscillators in the network, which become activated when an expected beat in the pattern fails to appear. The model makes predictions about the relative difficulty of patterns and the effect of deviations from periodicity in the input.

The Phenomenon

The nested periodicity that defines musical, and probably also linguistic, meter appears to be fundamental to the way in which people perceive and produce patterns in time. Meter by itself, however, is not sufficient to describe patterns which are interesting or memorable because of how they *deviate* from the metrical hierarchy. The simplest deviations are rests or gaps where one or more levels in the hierarchy would normally have a beat. When beats are removed at regular intervals which match the period of some level of the metrical hierarchy, we have what we will call a *simple rhythmic pattern*. Figure 1 shows an example of a simple rhythmic pattern. Below it is a grid representation of the meter which is behind the pattern.




Figure 1: A Simple Rhythmic Pattern

Simple rhythmic patterns, and their processing, are the concern of this paper. While there is not much research directly concerning such patterns, considerable research on the processing of periodic sequences of various types (e.g., Povel & Essens (1985)) does bear on them.

A complete model of the processing of these patterns would have at least the following properties. (1) It would be representationally adequate; that is, it would have the capacity to distinguish in its internal state the sorts of patterns which people distinguish. (2) At the perceptual level, it would exhibit the well-known constraints or tendencies concerning meter (Lerdahl & Jackendoff, 1983; Povel & Essens, 1985). For example, it would tend to assign prominences to beats which recur at regular intervals and would experience difficulty with patterns in which weak beats or rests appear where this principle would lead prominent beats to be expected. (3) It would have the capacity to use stored knowledge of simple rhythmic patterns to produce as well as recognize patterns. (4) On frequent exposure to a particular pattern, it would come to favor that pattern over others. Ultimately it would build a repertoire of familiar patterns to be applied to perception and production tasks. (5) It would generalize across different tempos; that is, stored knowledge associated with a pattern at a particular tempo could also be applied to patterns at nearby tempos. (6) It would be robust to temporal noise and other deviations from perfect periodicity.

From a strictly descriptive perspective, simple rhythmic patterns can be defined in terms of metrical structures from which beats are removed. But meter in the form of a static, spatially arranged hierarchy of beat sequences does not give any direct clues about how temporal patterns are perceived or produced online and may not even reflect their internal representation. Music perceivers must gradually accumulate the evidence for a particular meter as a pattern is presented, must deal with the problem of temporal short-term memory, and must cope with inevitable deviations from perfect meter. A metrical structure tells us little or nothing about how this is accomplished. What is needed is a continually-running system with a bias in favor of periodicity.

Oscillator Models of Meter Perception

Oscillator models of the processing of periodic patterns (Large & Kolen, 1994; McAuley, 1995; Miller, Scarborough, & Jones, 1992) take periodic behavior to be the default. Once activated, the primitive computational units in these models naturally expect events to recur at regular intervals. In adaptive oscillator models (Large & Kolen, 1994; McAuley, 1995), the oscillators adjust their phase angles and/or their frequencies in accordance with an input pattern. Adaptive oscillator models have the capacity to respond relatively smoothly to minor changes in tempo and random deviations from perfect periodicity and can generalize across different tempos. In oscillator network models (Miller et al., 1992), oscillators can excite or inhibit one another. Networks permit the inclusion of particular constraints on the sorts of periodicities which tend to co-occur or to fail to co-occur. If the oscillators are also coupled with each other—that is, if they are allowed to influence each other’s phase angle and/or frequency—those which are mutually supportive can “carry each other along,” as one or the other responds to the input.

The Model

Basic Architecture

The model we propose is a network of coupled, adaptive oscillators. It is a neural network in that it consists of a set of simple processing units which operate in parallel and respond to each other through weighted connections. Like other neural networks, its units may be activated directly by inputs. The activation of a unit is thus a measure of the extent to which the “knowledge” represented by that unit is characteristic of the current input. In our network there is a single input/output unit which activates the units in the network during perception and is activated by them during production. The units in this network, like those in architectures such as Hopfield networks and Boltzmann machines, are joined by connections with symmetric weights, and the activation update rule is a discrete approximation to a continuously-running dynamical system.

Unlike the units in most neural networks, however, those in this network are also oscillators, each of which is characterized by an instantaneous phase angle and frequency. Since the oscillators are coupled, they update not only their activations, but also their phase angle and their frequency, in response to other units.

The zero phase of an oscillator is the point at which it is “beating.” The output of an oscillator—that is, what is passed to the input/output unit during production—is the product of its activation and a periodic function of its phase which peaks at its zero phase, such as the cosine.

The details of activation and coupling, which are omitted in this short paper, depend on the oscillators. The input/output unit, which emits pulses, activates oscillators whose zero phases are near the pulses.

Oscillators in the network differ from each other in terms of their resting frequency and their initial phase angle. In a full-blown version of the model, the range of resting frequencies would span the space of frequencies which people can detect in periodic patterns.

Meter

As in the model of Miller et al. (1992), the recognition or production of a particular meter in the model takes the form of the activation of oscillators with frequencies which are integer multiples of one another, and the propensity for finding and remembering metrical structure is implemented in the form of strong excitatory connections between the oscillators whose frequencies are integer multiples of one another.

In our model the relationship between oscillators which have the potential to participate in a particular metrical structure goes further than this, however. Such oscillators are organized in *metrical clusters*. Within each cluster, the relationship between the frequencies and the phase angles of the oscillators is not allowed to vary. When one unit in a metrical cluster has its phase angle or frequency affected by the input unit or another oscillator outside the cluster, the change is propagated to all of the other units in the cluster so that their relative phases and frequencies remain constant. Since these adjustments are not done in parallel, the update cycle is randomized to avoid favoring certain units. Thus a cluster is something

Figure 2: Metrical Cluster designed for the pattern in Figure 1

Figure 2 shows a simple cluster which participates in patterns in 6/8 time at a particular tempo. It includes oscillators with two different periods, representing the dotted eighth note and the measure level in the meter. In the figure the size of the circles denoting the oscillators indicates their relative periods. There are six oscillators at the dotted eighth-note level, including three “rest oscillators” (explained below). Each of the three beat oscillators and rest oscillators is offset from the others by $1/3$ of its period, that is, by the duration of the eighth note. At the measure level, there are two oscillators, one which reaches its zero phase on the downbeat of each measure, the other which reaches its zero phase on the fourth beat of each measure.

Rests

To represent simple rhythmic patterns, we need a way to distinguish patterns with rests at different points within the meter. In the model, we treat rests as internal events comparable to beats; thus the recognition that a rest appears at a particular place in a pattern corresponds to the explicit activation of a rest unit, an oscillator whose zero phase is aligned with the position of the rest.

Unlike beat oscillators, rest oscillators are inhibited by the input unit and they do not affect the phase angle of the cluster. Rest oscillators are activated only by the beat oscillators in their cluster.

Pattern Reproduction

The excitatory and inhibitory connections shown in Figure 2 specify how the cluster would reproduce the pattern from Figure 1. Note that slower oscillators inhibit and excite faster ones. This influence is strongest when both units are near phase zero. Through this, each of the six faster oscillators can represent more than one beat or rest in the input pattern. For example, as one of the slower units approaches its phase zero, it may excite a faster oscillator; later another slower unit may come to phase zero and inhibit the same faster oscillator. When slower units exert this kind of excitation and inhibition on multiple beat and rest oscillators, a kind of chunking occurs such that a single slower unit is able to cause the reproduction of a whole series of input beats. This distribution of the input pattern among a relatively small number of oscillators suggests that a network of metrical clusters might be able to reproduce long patterns without an explosive increase in the number of units.

Figure 3 shows the summed output of this network following 80 presentations of the pattern. Though only positive-valued output is displayed, rest units in fact produce negative output when they pulse, serving to inhibit beats which occur at the same time. The summed output of the network is a measure of how well the network reproduces the input pattern. Note that the output pulses are stronger in places where the pattern repeats. This suggests that metrical clusters capture not only the actual beats and rests of a pattern but also the relative strengths of each event in the metrical structure.

This hand-wired network serves to show that the metrical cluster architecture is capable of representing simple rhythmic patterns. One long-term goal is to have the network learn the appropriate connections. Another goal is to explore more fully the capacity for chunking as described above.




Figure 3: Sum of output for Metrical Cluster in Figure 2

Conclusions

We have sketched a model in which the activation of oscillators in a network represents the perception or production of a beat or a rest and in which the hierarchical relations defining meter are represented in the form of hard-wired clusters of oscillators.

Clearly the model that we describe remains to be tested under various circumstances: in a larger network consisting of clusters with different fundamental frequencies and with input patterns which deviate from perfect periodicity. We must also test the learning algorithm which we are developing to demonstrate that the network can learn to favor certain rhythmic patterns over others.

However, the model already allows us to make certain predictions. When a unit in a metrical cluster adjusts its phase in accordance with the input, the other units in the cluster follow suit. This means, for example, that for the 6/8 pattern described above, an eighth note which appears late and causes the eighth-note oscillator to shift its phase angle back would also cause all of the other oscillators in the cluster to shift. That is, the next dotted eighth-note would be expected later as well. Furthermore, the model makes detailed predictions about what sequences of sub-patterns lead to simpler rhythmic patterns. In particular those patterns in which the subsequences make use of some of the same oscillators should be relatively easy to process.

References

- 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*. MIT Press, Cambridge, MA.
- McAuley, J. D. (1995). *On the Perception of Time as Phase: Toward an Adaptive-Oscillator Model of Rhythm*. Ph.D. thesis, Indiana University, Bloomington, IN.
- Miller, B. O., Scarborough, D. L., & Jones, J. A. (1992). On the perception of meter. In Balaban, M., Ebcioğlu, K., & Laske, O. (Eds.), *Understanding Music with AI: Perspectives on Music Cognition*, pp. 429–447. MIT Press, Cambridge, MA.
- Povel, D.-J. & Essens, P. (1985). Perception of temporal patterns. *Music Perception*, 2, 411–440.