Journal of Neuroscience Methods 183 (2009) 95–106

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

journal homepage: www.elsevier.com/locate/jneumeth

Pulse-coupled neuron models as investigative tools for musical consonance

B. Heffernan[a,c], A. Longtin[a,b,*]

[a] Centre for Neural Dynamics, University of Ottawa, Ottawa, Canada
[b] Department of Physics, University of Ottawa, 150 Louis Pasteur, Ottawa, Ontario K1N 6N5, Canada
[c] Systems Science Program, University of Ottawa, Canada

ARTICLE INFO

Article history:
Received 25 May 2009
Received in revised form 29 June 2009
Accepted 30 June 2009

Keywords:
Leaky integrate-and-fire
Mode locking
Mathematical models
Phase locking
Consonance
Dissonance
Noise
Coupled oscillators

ABSTRACT

We investigate the mode locking properties of simple dynamical models of pulse-coupled neurons to two tones, i.e., simple musical intervals. A recently proposed nonlinear synchronization theory of musical consonance links the subjective ranking from consonant to dissonant intervals to the universal ordering of robustness of mode locking ratios in forced nonlinear oscillators. The theory was illustrated using two leaky integrate-and-fire neuron models with mutual excitatory coupling, with each neuron firing at one of the two frequencies in the musical interval. We show that the ordering of mode locked states in such models is not universal, but depends on coupling strength. Further, unless the coupling is weak, the observed ratio of firing frequencies is higher than that of the input tones. We finally explore generic aspects of a possible synchronization theory by driving the model neurons with sinusoidal forcing, leading to down-converted, more realistic firing rates. This model exhibits one-to-one entrainment when the input frequencies are in simple ratios. We also consider the robustness to the presence of noise that is present in the neural firing activity. We briefly discuss agreements and discrepancies between predictions from this theory and physiological/psychophysical data, and suggest directions in which to develop this theory further.

1. Introduction

1.1. Theories of consonance

Despite centuries of theories and experiments, the precise neural basis for our perception of consonance and dissonance is still largely unknown. This is so for both the mechanisms involved and their location. Studies of consonance have focused almost exclusively on the perception of two simultaneous tones. This superposition can include only the two fundamental frequencies, or these frequencies plus their harmonics with a specified amplitude distribution for these harmonics. Each such sound is termed as a complex tone, and the presence of two tones, pure or complex, constitutes a dyad or a musical interval.

The problem of the mechanism underlying subjective ranking of an interval along the consonance to dissonance axis has received most attention at the psychophysical level, since experimental work involving the simultaneous processing of two periodic stimuli is challenging and limited by available recording technologies and suitable experimental preparations at the neuronal level. Yet the solution of this problem offers the exciting possibility of explaining a connection between simple subjective percepts and simple stimulus combinations in biophysical terms, and accordingly there are ongoing efforts to expose the neural basis of consonance evaluations.

In this work, we explore nonlinear dynamical models of neurons driven by periodic stimuli making up a musical interval, but also driven by each another through mutual excitatory coupling. The hope is to improve the biophysical realism of these models, explore the issues involved in mapping their modeled activity onto experimentally measurable activity, and reveal what aspects of synchronization—if any—may be at play.

The simplest and probably the oldest theory of consonance is that of Pythagoras. He observed that consonant mixtures of two tones occurred when the frequencies were in simple integer ratios. Helmholtz (1877) discussed consonance in the more general context of complex tones, which differ from pure tones in that they have power at harmonics of the fundamental. He proposed that dissonance is proportional to the number of frequency components present in the two complex tones that produce beats, i.e., whose frequency difference is within the so-called critical bandwidth (Kameoka and Kuriyagawa, 1969; Plomp and Levelt, 1965;

* Corresponding author at: Centre for Neural Dynamics, University of Ottawa, Ottawa, Canada. Tel.: +1-613-562-5800.
E-mail addresses: bheff028@uottawa.ca (B. Heffernan), alongtin@uottawa.ca (A. Longtin).

0165-0270/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jneumeth.2009.06.041

Roderer, 1995). For example, the sum of two such pure tone components of frequencies f_1 and f_2 and identical amplitude A can be written as:

$$A\sin 2\pi f_1 t + A\sin 2\pi f_2 t = 2A\cos\left(\frac{f_1 - f_2}{2}\right)\sin\left(\frac{f_1 + f_2}{2}\right) \quad (1)$$

When the frequencies are sufficiently close, this superposition produces a single perceived pitch at the average frequency $(f_1 + f_2)/2$ with a slowly varying intensity modulation known as the beat.

Another recent approach, based on timing nets, involves the analysis of population level distributions of all-order interspike intervals between firings (Cariani, 2001, 2004; Tramo et al., 2005). It relies on putative computations in the time domain with filters and coincidence detectors; the locus is not defined, but is thought to lie somewhere beyond the cochlear nucleus. It is based on the notion that harmonically related pitches share firing intervals at their common sub-harmonics. For example, in the case of a Perfect 5th, we have $2f_1 = 3f_2$, and thus $f_1/3 = f_2/2$ are common sub-harmonics of the two tones. The presence of these common sub-harmonics causes neural firings to be more correlated in time than for the case of non-harmonically related pitches. This in turn produces maximal pitch salience that can plausibly account for consonance of pairs of pure and complex tones.

Indeed, the simplicity of frequency ratios has played a central role in theories of intervallic consonance and dissonance. Sensory consonance is often distinguished from musical consonance. The former refers to consonance based on physical (i.e., acoustic) factors, and is, therefore, independent of musical conventions. Sensory consonance, which is considered to be a function of the aforementioned critical bandwidth, refers to the absence of amplitude fluctuations in two simultaneously sounded tones (because of their non-overlapping critical bands). Sensory dissonance refers to the "roughness" (very rapid amplitude fluctuations) that can result from simultaneously sounded tones with overlapping critical bands. By contrast, musical consonance is considered to result from tone compatibility, which is dependent on culture, convention, and context.

Moreover, musical consonance is applicable to sequential as well as simultaneous tones. From a psychoacoustic perspective, consonant intervals occur between 'compatible' tones and produce a 'feeling of stability', whereas dissonant intervals occur between 'incompatible' tones and cause instability (e.g., Aldwell and Schachter, 1989). Although the concepts of sensory and musical consonance differ, they are not completely independent (Bregman, 1990). For example, octaves have never been considered musically dissonant, and tones related by simple ratios, such as 2:3 and 1:2, are considered to be "stable" intervals across several musical cultures (Meyer, 1956).

Apart from issues of definition, there are outstanding problems in terms of the class of mechanisms that may underpin consonance ranking. As beating phenomena essentially arise from linear superposition of two sinusoidal waveforms, this concept of consonance and dissonance is a linear one. It has its limitations, which are nicely summarized in Shapira Lots and Stone (2008): consonance ratings can change beyond the critical bandwidth, can occur without the presence of harmonics, and cortical lesions reveal that there are specialized neuronal pathways dedicated to dissonance/consonance assessments (Peretz et al., 2001; Tramo et al., 2001). Sequential processing of tones also suggest that consonance does not rely as much on beats as on simple frequency ratios (Schellenberg and Trehub, 1994a,b), and EEG responses seem to imply that consonance ratings are formed by processing of pitch relationships in the auditory cortex (Itoh et al., 2003). So it is clear, given the range of subtly nuanced psychophysical phenomena and of outstanding problems, that much work is needed to link electrophysiological recordings, biophysical models and psychophysics. The work presented here offers one direction towards this goal, in the context of synchronization theory.

1.2. Possible contributions of nonlinearity

Shapira Lots and Stone (2008) recently proposed a synchronization theory of consonance that goes beyond the linear beating theory of Helmholtz. It is based on a striking observation made on numerical simulations of excitatorily pulse-coupled neuron models (see below): the progression from consonant to dissonant intervals is similar to the progression of step sizes on a 'Devil's Staircase'—the step sizes themselves being proportional to the width of Arnold tongues in nonlinear coupled oscillators. These technical concepts refer to the range of parameters over which a given mode locking of firing patterns is seen. A mode is defined as the frequency of one oscillator (e.g., a periodically firing neuron). Mode locking describes the phenomenon where the frequencies of two oscillators remain in a given ratio for some finite range of parameters. The fact that the oscillators adjust their frequency to maintain the same ratio is a sign of nonlinear synchronization.

For example, imagine neuron A firing at a fixed frequency. If it becomes excited periodically by neuron B, it may tend to synchronize its firings with that of neuron B. In the simplest case, there is a one-to-one correspondence between the firings of neurons A and B, i.e., one-to-one (1:1) mode locking. Alternately, for another parameter setting such as a lower amplitude of coupling, neuron A may fire only once for every two firings of neuron B, i.e., there is a 1:2 mode locking.

The general theory of nonlinear oscillators states that a ratio of $n + n' : m + m'$ can be found at parameters between those for which $n:m$ and $n':m'$ occur (the so-called Farey sequence—see, e.g., Hilborn (1994) for a general theory, and Glass and Mackey (1988) for specific applications to neuron models). There is in fact a universal sequence of mode locking ratios that appear as the ratio of the driving frequency (neuron B's frequency f_B) to the natural frequency (that of neuron A, i.e., f_A, in the absence of input from B) is increased—independently of the details of the models for the oscillators. Note here that f_B is not influenced by neuron A, i.e., the coupling is one-directional. One can then make a plot where the abscissa is the ratio of natural frequencies f_A/f_B, and the ordinate is the ratio of actual firing frequencies of the two coupled oscillators. Such a plot (examples are shown below) is known as a Devil's Staircase. This is called a staircase because it exhibits flat "steps" (actually, an infinite number of them) in which each corresponds to a mode locking. In other words, each step corresponds to a parameter range (e.g., a range of forcing frequencies) over which the same ratio of firing frequencies is seen at the output of the coupled oscillators. Strictly speaking the standard Devil's Staircase is defined for the so-called sine circle map (Glass and Mackey, 1988). In decreasing order of width, the steps correspond to 1:1, 1:2, 1:3 and 2:3 (same width), 2:5 and 3:5, etc., i.e., the steps decrease in width as higher integers occur in their fractional representation of the mode locking.

Conversely, the width of each step is also a measure of how robust a mode locking ratio is, i.e., of how easy it is to observe given variations in system parameters. The 1:1 step (unison in musical terms) is larger than the 2:1 step (octave), which is larger than the 3:2 step (Perfect 5th), larger than the 4:3 step (Perfect 4th) and so on. Shapira Lots and Stone (2008) observed that this sequence had notable similarity with the subjective ranking of consonance for musical dyads (see Table 1).

There have been recent dynamical approaches to perception, in which for example pitch perception relies on stable neural activity patterns known as dynamical attractors (Cartwright et al., 2001). The emphasis of such approaches is the nonlinearity of the complex neuronal systems at work. The perception of a beat frequency

Table 1
Mode stability as measured according to the authors' simulation results (column 4), and as measured by Shapira Lots and Stone (2008), for similar parameter values ($\alpha = 100$, $\varepsilon = 0.5$). Neither presents a strong correspondence with global subjective consonance rankings.

Interval	Ratio	Consonance ranking (Schwartz et al., 2003)	Ranking of mode stability ($\varepsilon = 0.5$)	Ranking of mode stability (Shapira Lots and Stone, 2008) ($\varepsilon = 5$)
Unison	1:1	1	**1**	**1**
Octave	1:2	2	3	**2**
Perfect 5th	2:3	3	2	**3**
Perfect 4th	3:4	4	**4**	**4**
Major 6th	3:5	5/6	**5**	**5/6/7**
Major 3rd	4:5	6/5	**6**	**5/6/7**
Minor 3rd	5:6	8	7	5/6/7
Minor 6th	5:8	7	10	8
Major 2nd	8:9	10/11	9	9
Major 7th	8:15	12	13	10
Minor 7th	9:16	11/10	12	**11**
Minor 2nd	15:16	13	11	(12)
Tritone	32:45	9	8[a]	(13)

[a] Here, we have measure the stability of the mode corresponding to 5:7, as the resulting difference in the second tone is imperceptible (less than 8 cents!)

requires nonlinearity in the auditory periphery (such as at the cochlea), since the beat frequency is not present in the linear superposition of two pure tones. The combination of sub-threshold activity with sufficient neuronal noise is thought to enable the detection of the missing fundamental through the so-called 'ghost stochastic resonance effect' (Chialvo et al., 2002; Chialvo, 2003). The detection by neurons of slow envelopes associated with narrowband signals relies on nonlinear thresholding with low noise (Middleton et al., 2006). The proposal by Shapira Lots and Stone (2008) follows this nonlinear trend, and is interesting given its fresh approach to the consonance problem.

1.3. Further explorations of a synchronization theory

Our study focuses on simple neuron models with excitatory pulse-coupling, as was used in the work of Shapira Lots and Stone (2008). In the discussion we comment on possible future extensions of our work to include inhibitory connections. The mode locking paradigm of Shapira Lots and Stone (2008) has not been tested against the recorded activity of nerve cells in the presence of dyads. Their scheme has only been illustrated numerically, assuming that each of two neurons fire, in the absence of coupling, at one of the frequencies present in the dyad. For example, for a Perfect 5th with a 256 Hz fundamental, the bias current of one neuron is adjusted until it fires periodically 256 times per second, and the other is adjusted to fire periodically 384 times per second (2:3). However, to our knowledge, neurons that fire periodically at the frequency of a tone stimulus have not been found in the cortex nor elsewhere (Tramo et al., 2005). In fact, cortical firing rates are typically low except when responding transiently to inputs. This model is thus seen more as a caricature of how nonlinearity and synchronization may arise in, e.g., auditory cortex, serving as a basis for exploring this theory of consonance.

Also, while it is clear that most neurons do exhibit a stochastic component to their firing, it is not clear what the balance of determinism and noise is needed to replicate the activity of cells involved in the ranking of consonance. There are numerous cell populations and intricate circuitry from the cochlear nucleus up to the auditory cortex (Tramo et al., 2005; Joris et al., 2003). The auditory afferents impinging on the cochlear nucleus already exhibit significant randomness in response to a pure tone. The firings are phase locked to the tone, but are separated by a random integer number of periods of the tone. This down-conversion of the input frequency to the output frequency is reproducible in terms of mathematical models that mix determinism and noise (see, e.g., Longtin, 1993 and references therein, and Cariani (2001) in the context of musical perception).

Other questions arise. How closely does the consonance ranking match psychophysical data for different coupling strengths and fundamental frequencies? How does it behave when neuronal noise is present? From the point of view of nonlinear dynamics, Devil's Staircases with strict universal properties are limited to weak coupling scenarios where one oscillator is driven by another (Glass and Mackey, 1988; Coombes and Bressloff, 1999). There is little known about the general properties of such "staircases" in the context of mutually coupled leaky integrate-and-fire neurons. Hereafter we nevertheless refer to the ensuing staircases as "Devil's Staircases" for simplicity. In fact, we notice in Shapira Lots and Stone (2008) that the stability measures are based on mode locking ratios that differ from the actual ratio of natural frequencies driving the neurons (see Section 3.2). For example, the step corresponding to the octave ratio (the neurons are mode locked 1:2) is seen when the input frequencies are in a ratio of around 1:3. Associating mode locking at one interval to the consonance of another interval undermines the theory. Further, there is little known about the synchronization properties of coupled oscillators, when each one is driven at substantially different frequencies (Pikovsky et al., 2003). Here we explore these properties further numerically in the context of consonance.

Finally we will consider an elaboration of the synchronization theory in which the neurons do not individually (i.e., without coupling) fire at the tone frequency. The higher processing pathways, in auditory and other similar pathways such as the electrosensory system (Berman and Maler, 1999), implement a kind of down-conversion of the frequency representation to bring it into line with the dynamical capabilities of the neurons. In fact much may be gained by looking at other senses, such as the electric sense and mechano-reception, that also deal with multiple inputs of a harmonic nature (Eggermont, 1990). While a full modeling of the neurons in the auditory pathway is beyond the scope of our work, it is possible to ask what the study of simple generic neuron models driven by sinusoids can bring to our understanding of consonance, on the road to more realistic models. Resulting observations and predictions could help refine experiments to validate this or other theories.

Section 2 exposes the methods used to explore how simple neuron models can serve as investigative tools for the study of consonance phenomena. Results on mode locking as a function of coupling strength and noise, along with more realistic sinusoidally forced models with lower rates are presented in Section 3. This section also includes an analysis of the correspondence between consonance rankings and mode locking ratios. A discussion of our results and outlook onto future investigations are the subjects of Section 4.

2. Methods

The leaky integrate-and-fire (LIF) model is a simple neuron model that retains the minimal ingredients of membrane dynamics, but whose behaviors nonetheless map onto many known properties of real neurons. They are sufficient to mimic basic sub-threshold properties of neurons. They incorporate supra-threshold spiking artificially: when the voltage reaches a fixed threshold, chosen equal to 1 in our work, a firing (or spike or action potential) is said to have occurred, and is represented graphically either by a point or by a vertical arrow on top of voltage time series plots (see below). At the next numerical integration time step, the voltage is reset to a value chosen here as 0. The threshold and reset voltages can be rescaled to realistic values for a given cell without qualitatively affecting results. For example the threshold can be set at −55 mV as the Na^+ activation voltage, and the reset can equal a resting potential of, e.g., −70 mV. The dynamics of the coupled LIF model can be written as:

$$\begin{aligned}\frac{dV_1}{dt} &= -\frac{V_1}{\tau_1} + I_1 + \varepsilon E_{2\to 1}(t)\\ \frac{dV_2}{dt} &= -\frac{V_2}{\tau_2} + I_2 + \varepsilon E_{1\to 2}(t)\end{aligned} \tag{2}$$

Here τ_1 and τ_2 are membrane time constants, chosen equal to 1 below. $E_{2\to 1}(t)$ represents the effect of neuron 2 (LIF2) on neuron 1 (LIF1), and vice versa for $E_{1\to 2}(t)$. The parameter ε represents the strength of coupling between the neurons. I_1 and I_2 represent the bias to the cells (in units of current divided by capacitance C, where C is set to 1). I_1 is chosen so that the neuron fires (reaches threshold) periodically at frequency $f_1 = 256$ Hz, the fundamental tone chosen for our study, and I_2 is chosen so that it fires at the other frequency f_2 in the interval. In general, the frequency of firing of an LIF model is related to the bias current I by the formula:

$$f^{-1} = \tau \ln \frac{I\tau}{I\tau - 1} \tag{3}$$

for $I\tau > 1$ which is the neural oscillator regime. This can be found simply by looking at the solution of a single LIF model (with $\varepsilon = 0$): $V(t) = I\tau(1 - e^{-t/\tau})$ and equating it to the threshold value after one period $T = f^{-1}$. When a neuron fires, an action potential is assumed to propagate to the other neuron, where it causes a synaptic current in the form of an alpha function with the time course:

$$E(t) = \alpha^2 t e^{-\alpha t}\Theta(t) \tag{4}$$

Here this formula represents what a spike at time zero contributes to the post-synaptic cell. $\Theta(t)$ is the Heaviside function which is 0 for $t < 0$ and 1 otherwise. The strength of this pulse-coupling between the oscillators is thus determined by $\varepsilon E(t)$. For numerical work involving many spikes, one has to sum many such alpha functions appropriately shifted in time to compute the ongoing effect of one cell on another. This involves keeping track of every firing time and numerically evaluating a continually growing list of alpha functions as the simulation proceeds. Further, the exponential evaluations are computationally costly, and amount to insignificant contributions after a few time constants. So instead, an equivalent procedure consists in modeling each synapse by two state variables:

$$\begin{aligned}\frac{dE_i}{dt} &= y_i\\ \frac{dy_i}{dt} &= -2\alpha y_i - \alpha^2 E_i + \sum_k \delta(t - t_{jk})\end{aligned} \tag{5}$$

Here the index $i = 1, 2, \ldots$ and j is the opposite of i, such that t_{jk} is the kth firing time of neuron j, and the sum is over all such firing times. The "delta" functions in the sum are Dirac delta functions, commonly used in computational neuroscience to mimic a spike arriving at a presynaptic terminal by focusing solely on its time of arrival. Numerical integration of the differential equations for each voltage and for the respective alpha functions of each neuron was done with an Euler scheme when no noise was present. The *input frequency ratio*, referred to below as *the intrinsic* or *natural frequency ratio*, was determined as the ratio of firing frequencies when the cells are uncoupled. The *output mode locked ratio* is the ratio of firing frequencies actually achieved in the coupled situation after transients have died out.

The above model, studied by Shapira Lots and Stone (2008), and also analytically earlier by Coombes and Lord (1997), will be studied below for different values of coupling. It will also be extended in two more realistic directions. To take noise into account, we will consider the effect of noise on the current–balance equation:

$$\begin{aligned}\frac{dV_1}{dt} &= -V_1 + I_1 + \varepsilon E_1(t) + \xi_1(t)\\ \frac{dV_2}{dt} &= -V_2 + I_2 + \varepsilon E_2(t) + \xi_2(t)\end{aligned} \tag{6}$$

The noises ξ_i are independent Gaussian white noises with zero mean. For simplicity both noises were given the same intensity D defining the autocorrelation function of the noise $\langle \xi(t)\xi(s)\rangle = 2D\delta(t - s)$. These equations were integrated using a standard Euler–Maruyama algorithm: the deterministic part of the dynamics is integrated with an Euler method, while the noise term at each time step contributes the value $N\sqrt{2D\Delta}$ where N is a random Gaussian number of mean zero and variance one, and Δ is the integration time step. Output mode locking ratios were determined by averaging the firing activity over long stretches of the numerical solutions; they are thus mean output mode locking ratios when noise is present.

The last model we consider is the sinusoidally forced, but noise-free, pulse-coupled LIF system:

$$\begin{aligned}\frac{dV_1}{dt} &= -V_1 + I_1 + A_1 \sin(2\pi f_1 t) + \varepsilon E_1(t)\\ \frac{dV_2}{dt} &= -V_2 + I_2 + A_2 \sin(2\pi f_2 t) + \varepsilon E_2(t)\end{aligned} \tag{7}$$

Without coupling ($\varepsilon = 0$) each neuron is known to exhibit mode locking to the periodic input (Keener et al., 1981). The frequencies of the input pure tones here are simply f_1 and f_2—there is no need for a calibration using Eq. (2). The threshold and reset are 1 and 0, respectively, for each cell as above. The chosen biases in this case are smaller than the values for Eq. (1) (see below). In fact, sub-threshold dynamics are used, such that firings cannot occur when coupling strength and forcing amplitudes A_1 and A_2 are set to zero. The amplitudes of the pure tones are set by A_1 and A_2. This formulation thus further allows an investigation of stimulus intensity by varying the amplitudes. It can also be used to mimic sensitivity of different neurons to different tone frequencies, i.e., to incorporate tonotopic receptive field properties. This is not explored here, as the amplitudes are set equal to one another, but our work sets the stage for these further explorations.

For these simulations, we also show the relative firing phases for each cell. Every time a neuron fires, its "phase" is reset to zero; this phase is then assumed to increase linearly in time until it reaches 2π at the time of its next firing (Pikovsky et al., 2003). Having assigned each neuron a phase, it then becomes possible to compute a relative phase as the phase at which neuron 1 (or 2) is when neuron 2 (or 1) fires. This representation simply illustrates the mode lockings.

3. Results

3.1. Coupling strength and mode locking

We first look at the effect of coupling strength in Eq. (1) on mode locking as well as on the ordering of the mode locking ratios. Fig. 1

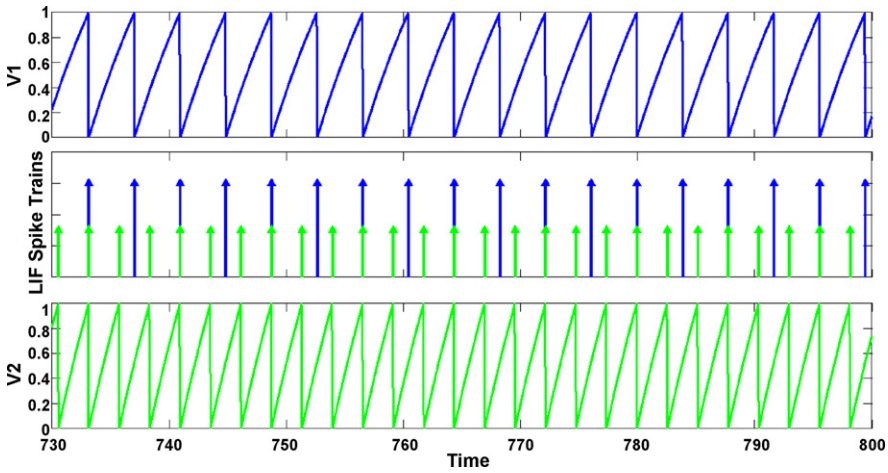

Fig. 1. The membrane potential of LIF1 (top panel, blue) and LIF2 (bottom panel, green), and their corresponding spike trains (middle panel). The bias currents in Eq. (1) were selected such that LIF1 fires at 256 Hz, and LIF2 fires at 384 Hz, constituting a Perfect 5th. Here, the LIFs are uncoupled ($\varepsilon = 0$) and thus, as expected, the resulting mode is just that of the natural firing ratio of 2:3, which is quite clear from the summed spike trains. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

shows the situation corresponding to a Perfect 5th interval (pure tone dyad) with a natural frequency ratio of 2:3. As the coupling strength is zero, the output mode locking ratio is the same as the natural ratio, as expected. Fig. 2 shows the mode locking effect resulting from an increase in coupling strength to $\varepsilon = 0.2$, resulting in a firing mode of 3:4. A coupling strength of $\varepsilon = 0.5$ results in a mode of 15:17 (not shown). In fact, the output mode locking ratio approaches 1:1 as ε increases.

It is also clear, by comparing the middle panels of Figs. 1 and 2, that the firing rates of both LIFs increase with increases in coupling strength. This is expected from Eq. (2) since the firing frequency is a monotonic increasing function of the bias current. The purely excitatory coupling causes a net positive contribution to the bias in each cell, which leads to a higher firing rate in each cell. This underlies the increase in the output mode locking ratio in comparison to the natural frequency ratio (Coombes and Lord, 1997).

3.2. Devil's Staircase

In order to reproduce the results of Shapira Lots and Stone (2008), a plot of the output mode locking ratio versus the natural frequency ratio was generated and is shown in Fig. 3. Here the natural frequency of LIF1 was set to 256 Hz, and the natural (intrinsic) frequency of LIF2 was swept from 230 to 1230 Hz in increments of 0.5 Hz by increasing the bias parameter I_2. A coupling strength of $\varepsilon = 0.8$ was used, and the resulting staircase structure is *nearly* indistinguishable from that published by those authors who used $\varepsilon = 8$. As they have confirmed this parameter choice (Shapira Lots and Stone, personal communication), it is not clear at this point why we must scale down their values of ε by a factor of 10 to produce a match to their results.

One can clearly see the largest step corresponding to the output ratio of 1 (1:1), followed by 0.5 (1:2), 0.66 (2:3), etc. The step size determines the robustness of a ratio, i.e., the range of system parameters, including natural frequencies, over which the ratio will be seen. The ordering of the step sizes and the correspondence with consonance rankings are discussed below using Table 1.

It is important to note here that all of the output mode locking ratios below 1 lie above the diagonal, and thus the mode ratios acquired with this coupling do not match the natural frequency ratios. This implies, for instance, that the octave—whose intrinsic frequency ratio is 1:2—produces a mode locking of roughly 16:25 (0.64), which is actually a fairly 'unstable' mode. For the parameter values used here, the coupled systems that actually synchronize to a mode locked ratio of 1:2 lie in the range of natural frequency ratios from 0.31 to 0.36 (roughly 1:3), which corresponds to a devi-

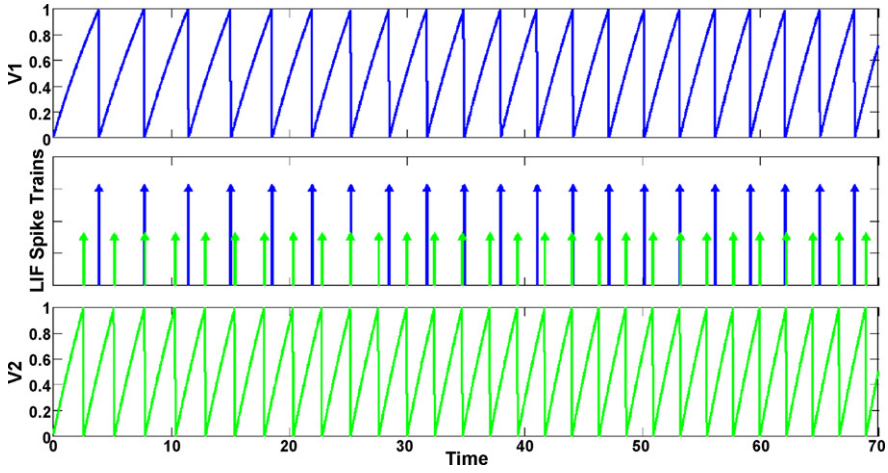

Fig. 2. A Perfect 5th with coupling strength $E = 0.2$, $\alpha = 0.9$. Here, the coupled LIFs, whose natural frequency ratio is 2:3, lock to a mode of 3:4. The frequency of discharge of each neuron has increased, as seen by comparing the middle panels of this panel with that of Fig. 1, since each neuron causes a net excitation in the other due to the excitatory coupling.

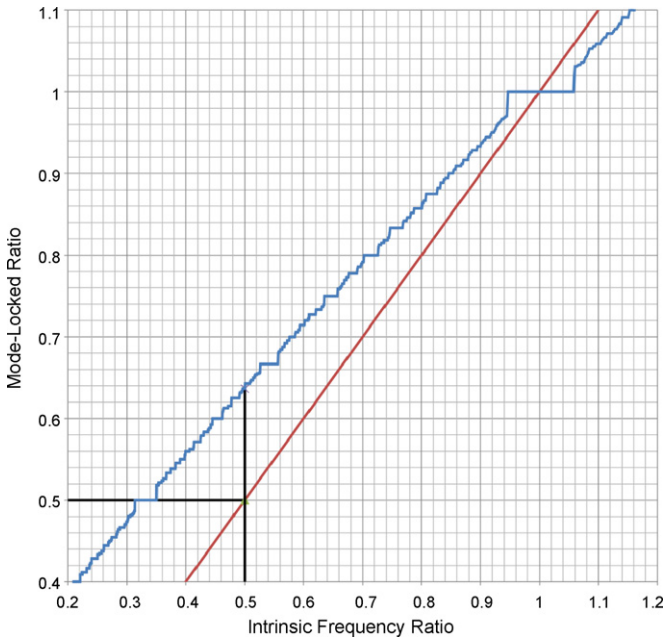

Fig. 3. 'Mutually Coupled-Excitatory Devil's Staircase' structure generated by holding LIF1 at 256 Hz and sweeping LIF2 from 230 to 1230 Hz. The coupling strength is $\varepsilon = 0.8$, $\alpha = 100$. The steps appear as lines due to the density of points occurring at the given modes. The red line indicates the diagonal. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

ation from the tonic (256 Hz) by an additional half-octave. This phenomenon is also observable in Shapira Lots and Stone (2008). It raises a concern regarding the strict association of step sizes with consonance ranking, as the ratio produced is shifted from the ratio actually desired. One must thus further explore the meaning of this association in this context.

3.3. Effect of noise on mode locking

Fig. 4 shows the effect of neuronal noise on the mode stability of the staircase, based on numerical simulations of Eq. (6). Here, due to the randomness of the solutions, the mode locking ratio is computed as a time average of the firing frequencies of both the oscillators. Clearly, for both coupling strengths shown, the noise has the effect of washing out the stair-like structure seen in Fig. 3. Only the larger steps remain visible for a given noise level. Noise does not change the slope of the curve, nor does it appear to change the relative sizes (and thus the size ordering) of the steps. A similar smearing of a Devil's Staircase by noise has also been reported in the context of electrosensory receptors in Chacron et al. (2000), which are very similar to mammalian primary auditory afferents (Carr, 2004). Interestingly, the staircase moves towards the diagonal as the coupling strength is reduced, i.e., the mode locking ratio acquired through coupling is closer to the natural ratio. Nevertheless, these results raise the question of how robust the synchronization for different ratios remains in the face of noise, an issue further studies will have to contend with. Perhaps other mechanisms are at work to enhance the mode lockings, such as coupling between multiple oscillators.

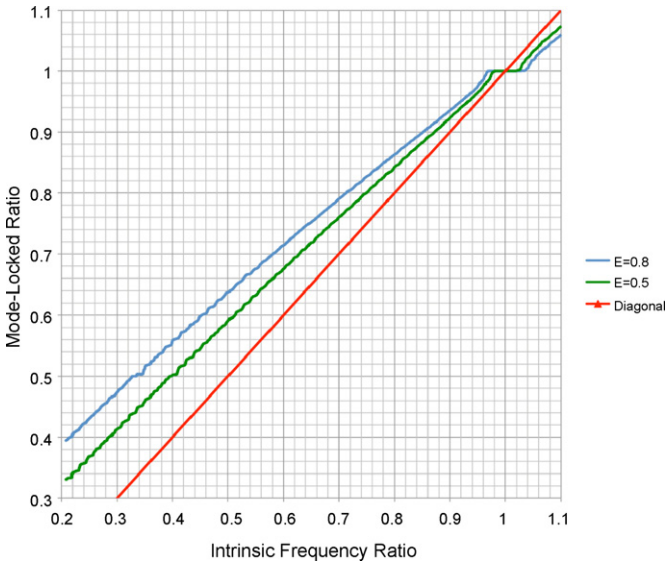

Fig. 4. Devil's Staircase generated in the same manner as in Fig. 3, here with noise added ($D = 0.1$), and for coupling strengths of $\varepsilon = 0.8$ (blue), and $\varepsilon = 0.5$ (red). Notice that the stairs (modes) of the staircase are almost entirely washed out by the noise. Both lie above the diagonal line shown here in green for all mode locked ratios less than 1, as expected. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

3.4. An analysis of correspondence of rankings

Again, under the assumption that Shapira Lots and Stone indicate ε values corresponding to a ten-fold increase to ours, we generated a staircase plot in an identical fashion to that mentioned above, now using $\varepsilon = 0.5$. Our assumption was further validated by a similarity in stability values, as measured by the width of each stair, i.e., the range of intrinsic frequency ratios over which a given mode locked ratio is maintained (see Appendix A for numerical values).

Table 1 shows the ranking of mode stability resulting from this frequency sweep at lower coupling placed beside those indicated by Shapira Lots and Stone. Bold values indicate a match in ranking between mode stability, and the mean value of consonance assessments due to Schwartz et al. (2003) derived by normalizing and averaging the rankings of several such studies (here '/' denotes equivalence).

From the table, it may seem that our reproduction of Shapira Lots and Stones' results is not such a faithful one. However, it is worth noting that the rank ordering of mode stability is parameter dependent, and thus the simple model Eq. (1) does not actually yield an invariant relation to consonance orderings as implied in Shapira Lots and Stone (2008). This indicates that this is not a true Devil's Staircase, since the width of the stairs do not map in a one-to-one correspondence with the simplicity of ratios regardless of model parameters.

Indeed, the notion of putting into correspondence the stability of mode locked ratios with dyadic tone ratios is further compromised by the fact that the 'staircase' has an average slope greater than one for all positive couplings. The slope tends to one as the coupling strength decreases (see Fig. 5 for a small coupling strength $\varepsilon = 0.08$). It can only be made to lie perfectly on the diagonal for $\varepsilon = 0$ (see also Fig. 5). In this latter case all modes reduce to points, yielding a stability measure of zero, and the staircase becomes a slide. This in itself does not exclude the possibility of 'Western' dyads locking to their identical mode locked ratios—even though the steps are now very small in comparison to previous results with stronger coupling. Nor does it exclude the possibility of the acquired mode's stability measures being ordered according to the simplicity of dyad ratios. But it does seem to imply that this perhaps overly simplistic model will never yield a perfect correspondence between the simplicity of frequency ratios and the stability of mode locked ratios in general. Given the similarity of some orderings however, it is nevertheless an interesting model from which to launch further explorations.

Table 2 shows the stability ranking of the output mode locking ratio acquired by the coupled LIF neurons driven at the dyadic ratio indicated, alongside the ranking of the stability of the mode of that same ratio. The stability of the actual mode acquired does not cor-

Table 2
The fourth column shows the rank ordering that results from measuring the stability of the *actual* mode acquired by the coupled LIFs, compared to Shapira Lots and Stone's measurement of the stability of the modes corresponding to the input ratio of the intrinsic frequencies, regardless of the fact that the LIFs *do not necessarily lock on to this ratio*. Clearly, correspondence between mode stability and consonance ranking is weak for both simulation results, as indicated by bold.

Interval	Ratio	Consonance ranking (Schwartz et al., 2003)	Stability ranking of actual mode acquired (measured from Fig. 3) ($\varepsilon = 0.8$)	Stability ranking of the mode corresponding to input ratio (measured from Fig. 3) ($\varepsilon = 0.8$)
Unison	1:1	1	**1**	**1**
Octave	1:2	2	9	**2**
Perfect 5th	2:3	3	11	**3**
Perfect 4th	3:4	4	3	5
Major 6th	3:5	5/6	7	7
Major 3rd	4:5	6/5	4	4
Minor 3rd	5:6	8	**8**	6
Minor 6th	5:8	7	5	8
Major 2nd	8:9	10/11	6	12
Major 7th	8:15	12	2	10
Minor 7th	9:16	11/10	**10**	**11**
Minor 2nd	15:16	13	**13**	**13**
Tritone	32:45	9	12	**9**

respond well with either the ranking of the stability of the mode locked ratios, nor with the global subjective consonance rankings. It is not much worse however than the stability of the mode locking ratio corresponding to the interval ratio.

3.5. Sinusoidally forced model

3.5.1. Mode locking

To further enhance the realism of our model in the context of synchronization, we performed a similar analysis to that presented above, this time using coupled LIFs with identical sub-threshold bias currents, each presented with a sinusoidal forcing at a frequency (pure tone) corresponding to one of the frequencies in the selected dyad (again using a fixed base tone of 256 Hz). The model dynamics are given by Eq. (7). Fig. 6 shows the response of the LIFs to a pure tone dyad stimulus corresponding to a Perfect 5th, here uncoupled to illustrate the acquisition of a mode of 2:3, as expected. A short transient phase appears characteristic of such sinusoidally forced sub-threshold LIFs over the parameters of interest here.

The effect of increasing the coupling strength ε while holding all other parameters fixed results in increasing (technically, non-decreasing) firing rates of the LIFs (Figs. 6–8). This results in a clear change of resulting mode, as evidenced by the relative phase of firing seen in the lower panels of Figs. 6–8.

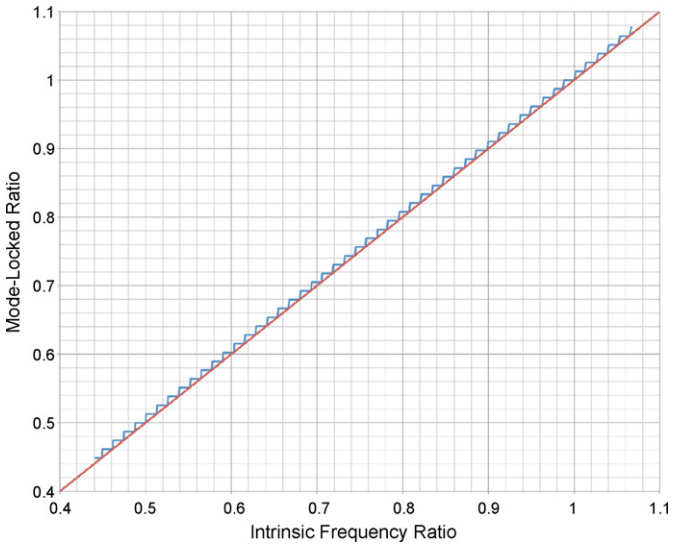

Fig. 5. Devil's Staircase with very low coupling ($\varepsilon = 0.08$). Although nicely aligned with the diagonal, the modes are of nearly uniform width. All dyads are 'perceptually stable', in the sense that they all lock to modes of significant width.

3.5.2. Analysis of the behavior of the model with sine forcing

Fig. 9 demonstrates the ability of a sub-threshold model neuron on its own (i.e., without coupling to another neuron) to either up-convert or down-convert the frequency of a sinusoidal stimulus. By fixing the sinusoidal amplitude to $A = 0.2$ and sweeping the stimulus frequency from 250 to 550 Hz, we see that the average firing rate follows a non-monotonic function. This is the case for the various bias currents we have investigated. What we are seeing here are mode lockings between the frequency of the input and the firing frequency of the neuron—although the input frequency is not influenced by the neuron firing (i.e., the coupling is uni-directional as opposed to bi-directional as has been studied up to now). Notice that the lower the bias current, the smaller the frequency range over which spikes are induced by the signal. This is a consequence of the low-pass filtering characteristic of the LIF model.

Notice also that for a bias of 0.96 (blue), the octave tone of 512 Hz is down-converted by a factor of 2. Down-conversion from tone frequency to neural firing frequency is already implemented by the primary auditory fibers, as discussed in the motivation leading to this model. Other conversions along the auditory pathway, which have not been fully elucidated, also occur to finally determine how primary auditory neurons fire according to a tonotopic representation (Joris et al., 2003; Tramo et al., 2005). The results here reveal that nonlinear mode locking effects may play a part in the down-conversion of frequency. This may be significant, as humans are capable of detecting pitches of stimuli whose frequencies are significantly greater than the upper limit of neuronal firing frequencies.

Fig. 10 shows the average firing frequency as a function of bias current for selected stimuli. It can be easily seen that for fixed amplitude of the stimulus, the LIFs will never generate spikes if the bias is too small. The low-pass filtering of the LIF is again shown clearly here, since as the stimulus frequency increases, the minimal bias *required in order to* make the LIF fire also increases. It is intriguing to see that the locking of the octave tone of 512 Hz (cyan) to the same average firing rate resulting from the tonic tone of 256 Hz (royal blue) occurs for a bias current of roughly *0.97*—which may relate to pitch invariance at octaves.

Finally, Fig. 11 presents an example of Devil's Staircase for this system Eq. (7) with coupling. Mode lockings are clearly seen in the numerous plateaus present. However the staircase is not monotonic as we have seen up to now. The structure of the mode lockings is quite complex, and its full analysis will be left for future work. So there is no apparent clear association between the size of steps

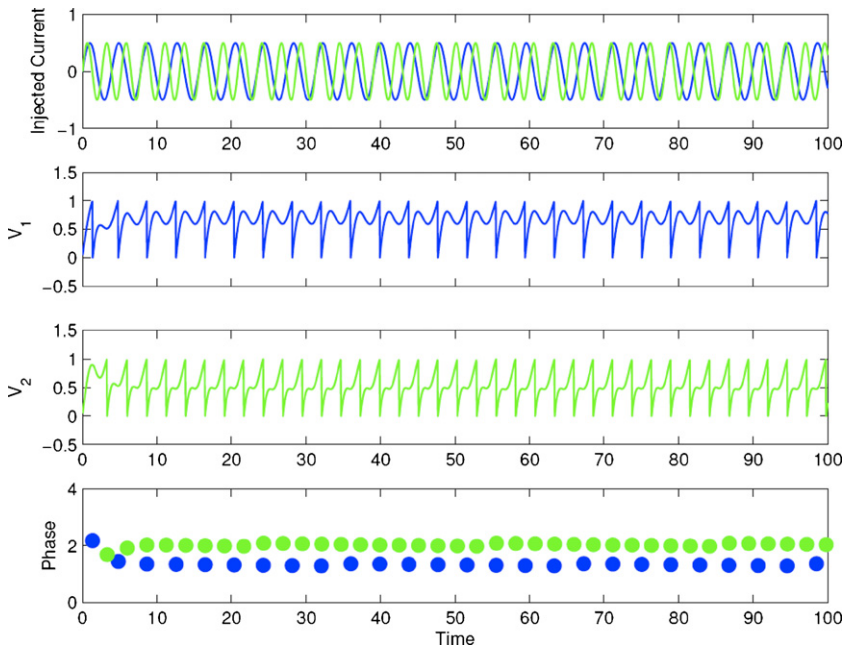

Fig. 6. (a) The sinusoidal current 'injected' into LIF1 (blue) and LIF2 (green), here representing a pure tone dyad stimulus of a Perfect 5th (256 and 384 Hz) with an amplitude of $A = 0.5$. (b and c) The membrane potential of LIF1 and LIF2 respectively, here uncoupled ($\varepsilon = 0$), each with a bias current of 0.93 and $\alpha = 0.9$. Notice the effect of the sinusoidal forcing on the shape of the membrane potential in comparison to the internally driven LIFs of Section 3.1. (d) The relative phase of the neurons shows clearly a mode of 2:3, as expected. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

and the simplicity of the ratio of the natural frequencies f_1/f_2. One feature stands out however: near the integer multiples of the tonic (so ratios of 1, 0.5, 0.33) the output mode locking ratio is around 1. This may again help to explain the pitch invariance experienced at octaves by humans and some primates.

4. Discussion

4.1. Effects of coupling and noise on ordering of consonance ratios

We have investigated simple models of nonlinear neural synchronization as a basis for ranking of musical intervals (dyads) along the consonance to dissonance axis. We first explored their very simple model in which each noise-free LIF neuron on its own fires at one of the two frequencies in the dyad. These neurons had mutual excitation via alpha functions triggered by each others firings. We obtained a plot of the ratio of firing frequencies with coupling to that without coupling, an extension of the standard Devil's Staircase used when one oscillator is driven uni-directionally by an external periodic rhythm. This plot agreed with that found in Shapira Lots and Stone (2008), although we had to use a coupling value ten times smaller then theirs. However other interesting features stood out. We noted that the position of the steps deviated significantly from the equivalent input ratio, especially for low to mid-range

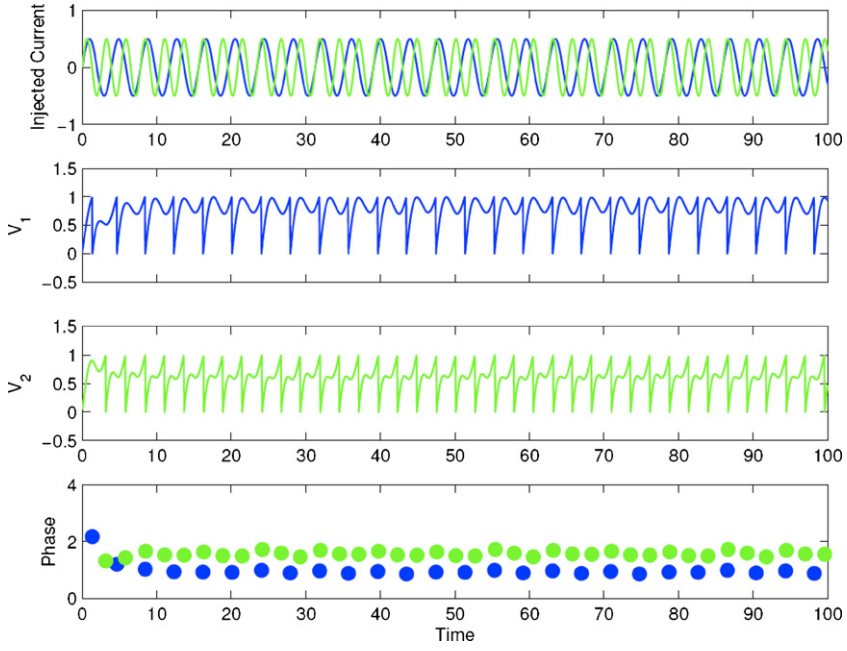

Fig. 7. Voltage time series as in Fig. 6 (Perfect 5th) but now with a coupling strength of $\varepsilon = 0.3$. Here, it is quite easy to see a mode of 2:3 occurring in the lower panel.

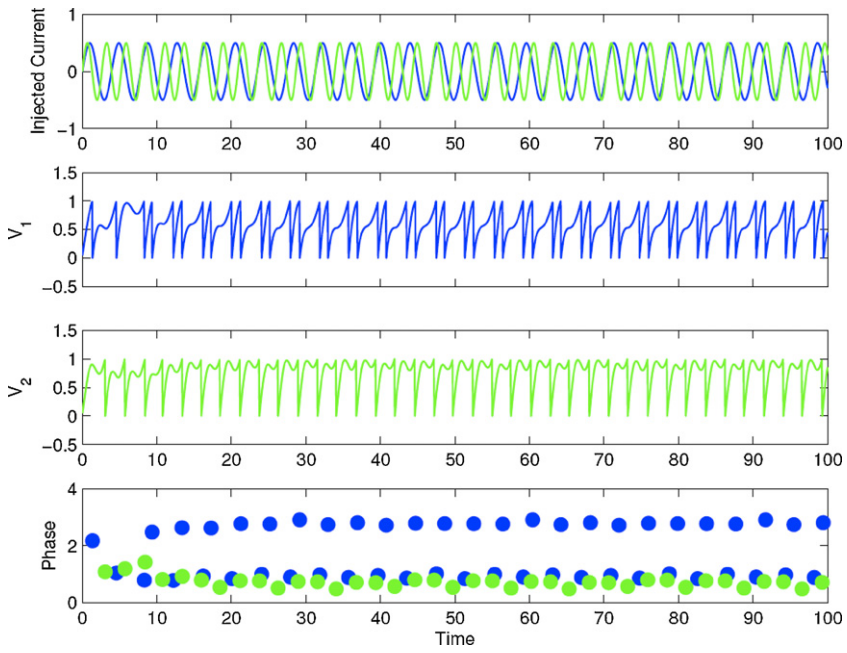

Fig. 8. Voltage time series as in Fig. 7, but here with the coupling strength increased to $\varepsilon = 0.5$. Notice the emergence of a new firing pattern for LIF1 (and for LIF2 as well, although more visually subtle). The resulting mode is significantly different than that seen in Fig. 7 for smaller coupling.

input ratios. This deviation pretty much disappeared as this ratio reached 1, but continues on above 1 (in fact, the value of 1 seems to be a pivot point around which the Devil's Staircases rotate as coupling strength varies). In other words, the plot deviated significantly from the diagonal. This means for example that input tones that are, e.g., in a 2:3 pattern (0.66 frequency ratio) produce a mode locking pattern that is not 2:3, but rather corresponds to a higher ratio.

Such deviations are intuitively expected. The excitatory coupling produces a time-varying input, the mean of which acts as an increased bias. Therefore, with coupling each neuron fires faster than without coupling. This nevertheless means that a consonance theory based on these simple neuron models must be re-interpreted to address this mismatch. This is especially so for the precise significance of the match between the orderings of the step sizes (robustness of mode locking ratios) and of the consonance rankings, since a given dyad (apart from unison) does not *generally* cause locking at the same ratio.

Our work further revealed that staircases are also seen for lower coupling strengths, with similar but not identical ordering of phase locking ratios to that seen for larger coupling (Tables 1 and 2). It also becomes more difficult to precisely estimate the size of the steps as the coupling strength is reduced, as the staircase acquires many steps of similar sizes, with all steps tending to lie on the diagonal.

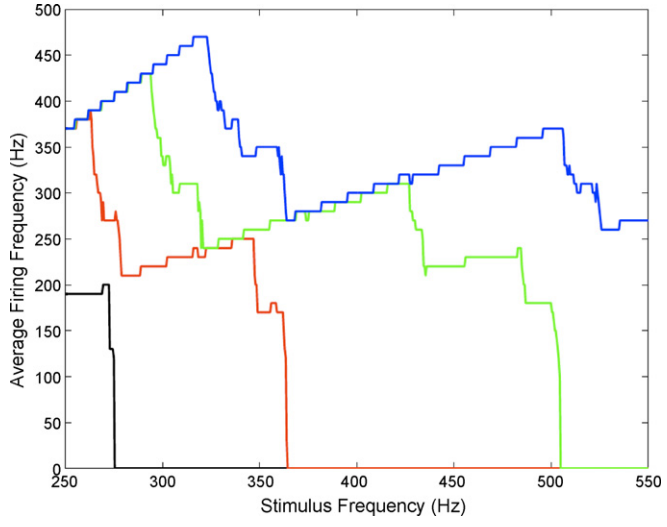

Fig. 9. The average firing frequency of a single LIF in Eq. (7) (uncoupled to the other LIF) as a function of stimulus frequency is plotted for a fixed bias current value (black = 0.9, red = 0.92, green = 0.94, blue = 0.96), for a constant stimulus amplitude of $A = 0.2$. Notice that only the highest bias current of 0.96 (blue) is capable of inducing spikes in the LIF neuron model across the entire octave spanned from 256 to 512 Hz. Here, the LIF being externally forced by a stimulus frequency of 512 Hz generates spikes at 256 Hz, down-converting the signal by a factor of 2. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

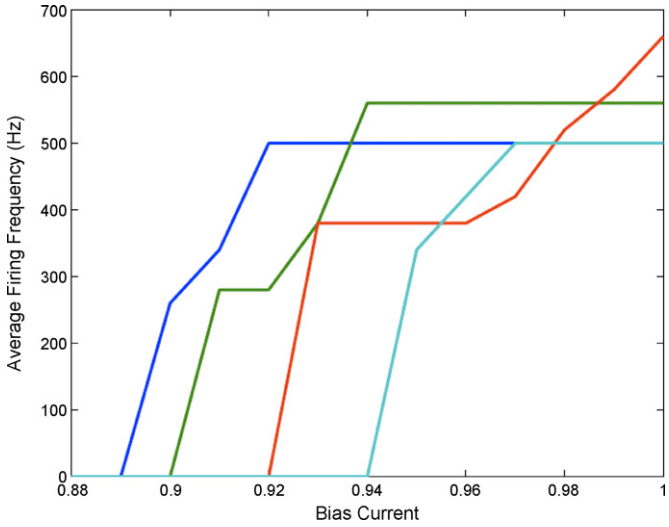

Fig. 10. Average firing frequency of a single uncoupled neuron in Eq. (7) as a function of bias current for stimulus tones of amplitude $A_1 = A_2 = 0.2$ presented at 256 (royal blue), 288 (green), 384 (red), and 512 Hz (turquoise). For the given amplitude, none of the stimulus tones bring the LIF neuron model to threshold for bias current values below 0.89. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

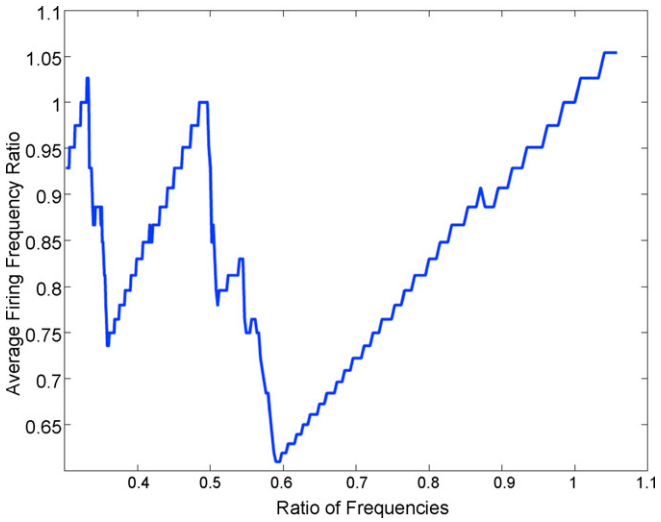

Fig. 11. Devil's Staircase for sinusoidally forced LIFs in Eq. (7) ($A_1 = A_2 = 0.4$, $\varepsilon = 0.4$, $I_1 = I_2 = 0.98$). Notice that the first and second octaves clearly mode lock to 1:1.

Again this is intuitively correct, since in the limit of no coupling, the ratio in equals the ratio out. The fact that the ordering varies with coupling strength suggests that it is not universal. It may be that the coupling present in the nervous system is such that it does have some degree of similarity with the subjective scale. This raises the issue of why such a coupling would be chosen.

Fig. 5 showed that, for very small ε values (here $\varepsilon = 0.08$), the mode locked ratios lie very close to the diagonal. This result should not be immediately dismissed due to a lack of a meaningful ordering of mode stability with regard to consonance rankings. Throughout human history, the experience of dyads has almost exclusively been of those in their complex form (until very recently, through the advent of sound cards, synthesizers, etc.). If we had never heard complex tones, we might not have any preferential regard for any dyad whatsoever, since, within a certain range, they may all be of nearly identical stability (as per Fig. 5). However, since this is clearly not the case, and we always experience a roughness proportional to the 'complexity' of the dyad (or its associated ratio), this assessment of 'roughness' may have become engrained or hard-wired. That is, from infancy, we begin to associate this roughness with the dyads, and thus our perceptual (and cultural) preference develops as a result. And pure toned dyadic experiments using left-ear right-ear separation of tones may still result in the same preferential rankings as a result of hard-wired responses.

Some of the ratios found in the model are perceptually equivalent, in that the acoustic stimuli they generate are not distinguishable to the average listener. This emphasizes that numerous ratios of varying complexity can be used to describe the same tonal phenomena, suggesting that there is a priori no special part played by simple ratios. It may also be that a relabeling according to the simplest ratio within the range generating a single percept will still result in a one-to-one correspondence with the global subjective rankings of consonance. For instance, relabeling the Tritone as 5:7 results in a better placement in terms of its actual consonance ranking. This results in a pitch difference of roughly 8 cents from the commonly employed ratio of 32:45, which is the limit of perception of the best professional piano tuners, and well inside the range of discrimination of the average ear.

4.2. Enhanced models

We then considered, for added biophysical plausibility, the effect of additive Gaussian white noise on this synchronization theory. The effect of the noise was to jitter the spike times and mode locked firing patterns, with the result that the overall shape of the firing mode lockings, as well as phase paths (not shown) was preserved. However, the staircase had less visible fine structure and more "rounded" steps. The orderings based on the steps that had measurable widths remained as in the noiseless case, as observed in externally forced neural models (Chacron et al., 2000). It would appear however that, unless, e.g., yet-to-be-modeled network effects come to the rescue, or noise-induced firings from sub-threshold dynamics play an important role, the theory could quickly lose its firm footing when the noise becomes moderate and the steps are washed away. It remains to be seen what noise level is actually at work in the cells responsible for consonance perception. Perhaps the averaged mode locking ratios play a significant role in such perception as they do in determining, e.g., the shape of tuning curves (Longtin, 2001).

Finally we have considered another novel modeling step that relaxes the rather unphysical requirement that each model neuron fire at one of the frequencies in the interval. The model LIFs are now driven by sinusoidal signals of adjustable frequency and amplitude. By adjusting the mean bias to the model cells, it is possible to choose the mean output firing rate of an LIF for a given forcing. This can then be calibrated against neural data in a structure that is putatively involved in consonance or dissonance ranking. This frequency can be made lower than the driving frequency, thus implementing a simple down-conversion of firing rates in this system (which would add to earlier effects such as stochastic phase locked firing in the cochlear nerve).

Such a model LIF neuron with sinusoidal fording is also known to exhibit phase locking to one sinusoidal input, with the usual universal orderings that characterize nonlinear systems (Keener et al., 1981)—i.e., there is already a Devil's Staircase at the single neuron level, prior to the coupling. Consequently, plots of firing frequency versus bias current are non-monotonic and exhibit much structure. Plots of firing frequency versus tone frequency nevertheless decay monotonically, a consequence of the fact that the LIF acts as a low-pass filter.

Pulse-coupling two such LIFs with excitatory synapses produces dynamical effects that will require much effort to analyze in detail. Our more limited goal here is to see whether there is any semblance of mode locking as in the Stone and Shapira Lots model and/or to consonance rankings. The simple answer is that there isn't any simple semblance. Several difficulties arise due to the mode locking that follow multiple complex staircases. However, one order seems apparent: for simple integer ratios of stimulus frequencies, the two LIFs are close to 1:1, i.e., the oscillators entrain one another (Fig. 11). Such entrainment may play a role in the perceptual effect of pitch invariance at octaves by signaling the presence of a simple input ratio. It may be a neural state of synchrony that is associated with the subjective experience of consonance. It remains to be seen what feature of this entrainment, such as the width of the peaks in Fig. 11, might be associated with the ranking of consonance.

4.3. Model extensions

We have chosen equal amplitudes for the pure tones in Eq. (7). This formulation allows an investigation of stimulus intensity by varying the amplitudes in a manner that relates, e.g., nonlinearly, to that of the actual acoustic intensities. These elaborations will be explored elsewhere, building on the equi-amplitude picture discussed here. The effect of noise on the pulse-coupled LIF system with sinusoidal forcing will also be of interest. The program of fully analyzing Eq. (7), to look at effects of frequency, amplitude, membrane and coupling time constants, etc., is a hefty one, and should begin with a proper non-dimensionalization to reduce the number of parameters.

Table A1
The actual mode stability measurements used to produce the rankings shown in Tables 1 and 2 computed from our simulation results. The final column to the right shows the measurements indicated in Shapira Lots and Stone (2008), which were used to produce the associated ranking shown in Table 1.

Interval	Interval ratio as indicated in Shapira Lots and Stone (2008)	Consonance ranking (Schwartz et al., 2003)	Stability of actual mode acquired ($\varepsilon = 0.8$)	Stability of mode corresponding to ratio of the interval ($\varepsilon = 0.8$)	Stability of mode corresponding to ratio of interval ($\varepsilon = 0.5$)	Stability of mode corresponding to ratio of interval indicated in Shapira Lots and Stone (2008) ($\varepsilon = 5$)
Unison	1:1	1	0.11150	0.11150	0.08400	0.075
Octave	1:2	2	0.00388	0.03691	0.02505	0.023
Perfect 5th	2:3	3	0.00261	0.02917	0.02648	0.022
Perfect 4th	3:4	4	0.02126	0.02201	0.01488	0.012
Major 6th	3:5	5/6	0.00768	0.01687	0.01480	0.010
Major 3rd	4:5	6/5	0.01231	0.02391	0.01306	0.010
Minor 3rd	5:6	8	0.00408	0.02126	0.00970	0.010
Minor 6th	5:8	7	0.00916	0.01236	0.00398	0.007
Major 2nd	8:9	10/11	0.00773	0.00408	0.00573	0.006
Major 7th	8:15	12	0.02917	0.00724	0.00074	0.005
Minor 7th	9:16	11/10	0.00308	0.00657	0.00086	0.003
Minor 2nd	15:16	13	0.00172	0.00159	0.00164	–
Tritone	32:45	9	0.00239	0.00768	0.00906	–

It will also be interesting to extend the models presented here to even more realistic neural populations which somehow share the information about the complex tones. A specific possibility is to assume that one neuron (or neural sub-population) is being driven by one tone (or complex tone) at one amplitude, but also by the other tone (or complex tone) at a different amplitude. The issue here is one of spectral receptive field, in which the proximity of tones in frequency space will reflect the proximity of the neurons they respectively excite, according to the tonotopic map. Accordingly, if the tones are close (as in a Minor 2nd) the neurons being excited may have very similar input made up of both tones with approximately equal weight. On the other hand, for tones much further apart in frequency (such as for a Perfect 5th with a 3:2 ratio) one neuron will be driven more by one tone than the other, and vice versa for the other. It will thus be interesting to investigate the synchronization properties of coupled populations of such neurons with dyad-dependent mixing ratios of their inputs amplitudes. One can only speculate at this point on the resulting phase locking structure and its robustness to noise.

Note that, although we have not investigated this point here, the sinusoidally forced model formulation enables the explorations of loudness effects on consonance evaluation, by adjusting the amplitude of the sine waves in some proportion to the stimulus intensity. It may also be of interest to study how fast mode locking is established under different modeling scenarios and see if there is any qualitative agreement with psychophysical studies on the speed of consonance ranking. Preliminary results indicate that rapid locking occurs for Eq. (7). It will also be interesting to see what inhibitory connections add to the dynamical behaviors discussed here, as it is also known to cause mode locking. It may allow special forms of mode locking all the while keeping the firing frequencies close to their values in the absence of coupling. In other words, output ratios may be closer to the input ratio. Another avenue is to investigate the connection between mode locking and the pitch salience seen in timing net approaches that involve population level interval estimation and coincidence detectors (Cariani, 2001, 2004).

Finally, data on humans and primates need to be reconciled with the synchronization picture, as for any other theory of consonance. Firing rates are much lower than tone frequencies; there are many random firings; and dissonance has been claimed to be associated with cortical firing rates modulated at slow beat frequencies, rather than complex locking ratios (see, e.g., Fishman et al. (2001) for data on macaque monkeys and event-related potentials in humans). It will thus be important to investigate how the activity from neuron models can be brought into correspondence with multi-unit activity in cortex.

Acknowledgements

This research was supported by the Natural Sciences and Engineering Research Council of Canada. The authors acknowledge useful conversations with Andreas Daffertshofer and Len Maler and technical help from Jason Boulet.

Appendix A.

See Table A1.

References

Aldwell E, Schachter C. Harmony and Voice Leading. 2nd ed. San Diego, CA: Harcourt Brace Jovanovich; 1989.
Berman N, Maler L. Neural architecture of the electrosensory lateral line lobe: adaptations for coincidence detection, a sensory searchlight and frequency-dependent adaptive filtering. J Exp Biol 1999;202:1243–53.
Bregman AS. Auditory Scene Analysis: The Perceptual Organization of Sounds. Cambridge, MA: MIT Press; 1990.
Cariani P. Temporal codes, timing nets, and music perception. J New Music Res 2001;30(2):107–35.
Cariani P. A temporal model for pitch multiplicity and tonal consonance. In: Proceedings, int. conf. music perception & cognition (ICMPC); 2004.
Carr CE. Timing is everything: organization of timing circuits in auditory and electrical sensory systems. J Comp Neurol 2004;26:131–3.
Cartwright JH, Gonzales DL, Piro O. Pitch perception: a dynamical-systems perspective. PNAS 2001;98:4855–9.
Chacron M, Longtin A, St-Hilaire M, Maler L. Suprathreshold stochastic firing dynamics with memory in P-type electroreceptors. Phys Rev Lett 2000;85:1576–9.
Chialvo DR, Calvo O, Gonzalez DL, Piro O, Savino GV. Subharmonic stochastic synchronization and resonance in neuronal systems. Phys Rev E 2002;65:050902.
Chialvo DR. How we hear what is not there: a neural mechanism for the missing fundamental illusion. Chaos 2003;13:1226–30.
Coombes S, Lord GJ. Intrinsic modulation of pulse-coupled integrate-and-fire neurons. Phys Rev E 1997;56:5809–18.
Coombes S, Bressloff P. Mode locking and Arnold tongues in integrate-and-fire neural oscillators. Phys Rev E 1999;60:2086–96.
Eggermont JJ. The correlative brain. Berlin: Springer; 1990.
Fishman Y, Volkov I, Noh M, Garell P, Bakken H, Arezzo J, et al. Consonance and dissonance of musical chords: neural correlates in auditory cortex of monkeys and humans. J Neurophysiol 2001;86(6):2761–88.
Glass L, Mackey MC. From clocks to chaos. Princeton: Princeton University Press; 1988.
Helmholtz H. On the sensations of tone as a physiological basis for the theory of music. New York, NY: Dover Publications; 1877.
Hilborn RC. Chaos and nonlinear dynamics. Oxford: Oxford University Press; 1994.
Itoh K, Suwazono S, Nakada T. Cortical processing of musical consonance: an evoked potential study. NeuroReport 2003;14(18):2303–6.
Joris PX, Schreiner CE, Rees A. Neural processing of amplitude-modulated sounds. Physiol Rev 2003;84:541–77.
Kameoka A, Kuriyagawa M. Consonance theory I. Consonance of dyads. J Acoust Soc Am 1969;45:1451–9.
Keener JP, Hoppenstadt FC, Rinzel J. Integrate and fire models of nerve membrane response to oscillatory input. SIAM J Appl Math 1981;41:503–17.
Longtin A. Stochastic resonance in neuron models. J Stat Phys 1993;70:309–27.

Longtin A. Effect of noise on the tuning properties of excitable cells. Chaos Solitons Fractals 2001;11:1835–48.
Meyer LB. Emotion and Meaning in Music. Chicago, IL: University of Chicago Press; 1956.
Middleton JW, Longtin A, Benda J, Maler L. The cellular basis for parallel neural transmission of a high-frequency stimulus and its low-frequency envelope. Proc Natl Acad Sci USA 2006;103:14596–601.
Peretz I, Blood AJ, Penhune V, Zatorre R. Cortical deafness to dissonance. Brain 2001;124:928–40.
Pikovsky A, Rosenblum M, Kurths J, Synchronization. A universal concept in nonlinear sciences. Cambridge, UK: Cambridge University Press; 2003.
Plomp R, Levelt WJM. Tonal consonance and critical bandwidth. J Acoust Soc Am 1965;38:548–60.
Roderer JG. The physics and psychophysics of music: an introduction. Berlin: Springer Verlag; 1995.
Schellenberg E, Trehub S. Frequency ratios and the discrimination of pure tone sequences. Percept Psychophys 1994a;56:472–8.
Schellenberg E, Trehub S. Frequency ratios and the perception of tone patterns. Psychonom Bull Rev 1994b;1:191–201.
Schwartz D, Howe C, Purves D. The statistical structure of human speech sounds predicts musical universals. J Neurosci 2003;23(18):7160–8.
Shapira Lots I, Stone L. Perception of musical consonance and dissonance: an outcome of neural synchronization. J R Soc Interface 2008;5(29):1429–34.
Tramo MJ, Cariani PA, Delgutte B, Braida LD. Neurobiological foundations for the theory of harmony in Western tonal music. Ann N Y Acad Sci 2001;930:92–116.
Tramo MJ, Cariani PA, Koh CK, Makris N, Braida LD. Neurophysiology and neuroanatomy of pitch perception: auditory cortex. Ann N Y Acad Sci 2005;1060(1):148–74.