

---

# A Graphical Model for Chord Progressions

---

chord progressions, graphical models, musical structure, music generation

## Abstract

Chord progressions are the building blocks from which tonal music is constructed. Inferring chord progressions is thus an essential step towards modeling long term dependencies in music. In this paper, a distributed representation for chords is designed such that Euclidean distances roughly correspond to psychoacoustic similarities. Graphical models observing chord progressions are then compared in terms of conditional out-of-sample likelihood. Both perceptual and statistical evidence show that binary trees related to meter are well suited to capture chord dependencies.

## 1. Introduction

Probabilistic models for analysis and generation of polyphonic music would be useful in a broad range of applications, from contextual music generation to online music recommendation. However, modeling music in general involves long term dependencies in time series that have proved very difficult to capture with traditional statistical methods. (Note that the problem of long-term dependencies is not limited to music, nor to one particular probabilistic model; see Bengio et al. (1994)). This difficulty motivates our exploration of chord progressions. Chord progressions constitute a fixed, non-dynamic structure in time and thus can be used to aid in describing long-term musical structure.

Tonal music comprises most of the music that has been written since J.-S. Bach (including contemporary pop music.) One of the main features of tonal music is its organization around *chord progressions*. A chord is a group of three or more notes (generally five or less.) A chord progression is simply a sequence of chords. In general, the chord progression itself is not played directly in a given musical composition. Instead, notes comprising the current chord act as central polarities for the choice of notes at a given moment in a musical

piece. Given that a particular temporal region in a musical piece is associated with a certain chord, notes comprising that chord or sharing some harmonics with notes of that chord are more likely to be present. In typical tonal music, most chord progressions are repeated in a cyclic fashion as the piece unfolds, with each chord having in general a length equal to integer multiples of the shortest chord length.

The interaction between the notes that are actually played and the notes comprising the chord progression are related to the meter of the piece. Meter is the sense of strong and weak beats that arises from the interaction among hierarchical levels of sequences 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 (Handel, 1993). For instance, most contemporary pop songs are built on 4-beat meters. In such songs, the first and third beats are usually emphasized. In terms of melodic structure, this indicates that notes perceptually closer to the chord progression are more likely to be played on these beats while more “dissonant” notes can be played on weaker beats.

These observations motivate our attempt to directly model chord progressions. Because the space of sensible chord progressions is much more constrained than the space of sensible melodies, our model could form an important part of a system that analyzes or automatically generates melodies. As an example, consider blues music. Most blues compositions are variations of a basic *same* 12 bar chord progression<sup>1</sup>. Identification of that chord progression in a sequence would greatly contribute to genre recognition.

In this paper we present a graphical model that captures the chord structures in a given musical style using as evidence a limited amount of symbolic MIDI<sup>2</sup>

---

<sup>1</sup>In this paper, chord progressions are considered relative to the key of each song. Thus, transposition of a whole piece has no effect on our analysis.

<sup>2</sup>MIDI stands for Musical Instrument Digital Interface, an industry-standard interface used on electronic musical keyboards and PCs for computer control of musical instruments and devices. In our work, we only consider note

data. One advantage of graphical models is their flexibility, suggesting that our model could be used either as an analytical or generative tool to model chord progressions. Moreover, a model like ours can be integrated into a more complex probabilistic transcription model (Cemgil, 2004), genre classifier, or automatic composition (Eck & Schmidhuber, 2002).

Cemgil (2004) uses a somewhat complex graphical model that generates a mapping from audio to a piano-roll using a simple model for representing note transitions based on Markovian assumptions. This model takes as input audio data, without any form of pre-processing. While being very costly, this approach has the advantage of being completely data-dependent. However, strong Markovian assumptions are necessary in order to model the temporal dependencies between notes. Hence, a proper chord transition model could be appended to such a transcription model in order to improve polyphonic transcription performance.

Raphael and Stoddard (2003) use graphical models for labeling MIDI data with traditional Western chord symbols. In this work, a Markovian assumption is made such that each chord symbol depends only on the preceding one. This assumption seems sufficient to infer chord symbols, but we show in Section 4 that longer term dependencies are necessary to model chord progressions by themselves in a generative context, without regard to any form of analysis.

Lavrenko and Pickens (2003) offer a generative model of polyphonic music that employs Markov random fields. Though the model is not restricted to chord progressions, the dependencies it considers are much shorter than in the present work. Also, octave information is discarded, making the model unsuitable for modeling realistic chord voicings. For instance, low notes tend to have more salience in chords than high notes (Levine, 1990).

Allan and Williams (2005) designed a harmonization model for Bach chorales using Hidden Markov Models. A harmonization is a particular choice of notes given a sequence of chord labels. While generating excellent musical results, this model has to be provided sequences of chords as input, restricting its applicability in more general settings. Our work goes a step further by modeling directly chord progressions in an unsupervised manner. This allows our proposed model to be directly appended to any supervised model without the need for additional data labeling.

onsets and offsets in the MIDI signal.

## 2. Representation

The generalization performance of a generative model depends strongly on the chosen representation for observed chords. A good representation encapsulates some of the psychoacoustic similarities between chords. One possibility we chose not to consider was to represent directly some attributes of Western chord notation such as *minor*, *major*, *diminished*, etc. Though inferring these chord qualities could have aided in building a similarity measure between chords, we found it more convenient to build a more general representation directly tied to the acoustic properties of chords. Another possibility for describing chord similarities is set-class theory, a method that has been compared to perceived closeness (Kuusi, 2001) with some success. In this paper, we consider a simpler approach where each group of observed notes forming a chord are seen as a single timbre (Vassilakis, 1999) and we design a continuous distributed representation where close chords with respect to Euclidean distance tend to be similar to listeners.

More specifically, the frequency content of an idealized musical note  $i$  is composed of a fundamental frequency  $f_{0,i}$  and integer multiples of that frequency. The amplitude of the  $h$ -th harmonic  $f_{h,i} = hf_{1,i}$  of note  $i$  can be modeled with geometric decaying  $\rho^h$ , with  $0 < \rho < 1$  (Valimaki et al., 1996).

Consider the function

$$m(f) = 12(\log_2(f) - \log_2(8.1758))$$

that maps frequency  $f$  to MIDI note  $m(f)$ . Then, for a given chord, we associate to each MIDI note  $n$  a perceived loudness  $l(n)$  with

$$l(n) = \max(\{\rho^h | \text{round}(m(f_{h,i})) = n\} \cup \{0\}) \quad (1)$$

where the function `round` maps a real number to the nearest integer. The `max` function is used instead of a sum in order to account for the masking effect (Moore, 1982). The quantization given by the rounding function corresponds to the fact that most of the tonal music is composed using the well-tempered tuning. For instance, the 3rd harmonic  $f_{3,i}$  corresponds to a note which is located one perfect fifth (i.e. 7 semi-tones) over the note corresponding to the fundamental frequency. Building the whole set of possible notes from that principle leads to a system where flat and sharp notes are not the same, which was found to be impractical by musical instrument designers in the baroque era. Since then, musicians used a compromise called the well-tempered scale, where each semi-tones are separated by an equal ratio of frequencies. Hence, the rounding function in Equation (1) provides


a frequency quantization that corresponds to what an average contemporary music listener experiences on a regular basis.

For each chord, we then have a distributed representation  $\mathbf{l} = \{l(n_1), \dots, l(n_d)\}$  corresponding to the perceived strength of the harmonics related to every note  $n_i$  of the well-tempered scale. Using octave invariance, we can go further and define a chord representation  $\mathbf{v} = \{v(0), \dots, v(11)\}$  where

$$v(i) = \sum_{n_j: 1 \leq j \leq d, (n_j \bmod 12) = i} l(n_j). \quad (2)$$

This representation gives a measure of the relative strength of each pitch-class<sup>3</sup> in a given chord. For instance, value  $v(0)$  is associated to pitch-class  $c$ , value  $v(1)$  to pitch-class  $c$  sharp, and so on. We see in Figure 1 that this representation gives similar results for two different voicings of the C major chord, as defined in Levine (1990).

Throughout this paper, we define chords by giving the pitch-class letter, sometimes followed by symbol # (sharp) to raise a given pitch-class by one semi-tone. Finally, each pitch-class is followed by a digit representing the actual octave where the note is played. For instance, the symbol  $c1e2a\#2d3$  stands for a 4-note chord



with a  $c$  on the first octave, an  $e$  and an  $a$  sharp ( $b$  flat) on the second octave, and finally a  $d$  on the third octave.

We have computed Euclidean distances between chords represented in the distributed chord space given by Equation (2) and found that they roughly correspond to perceptual closeness, as shown in Table 1. Each column gives Euclidean distances between the chord in the first row and some other chords that are represented as described here. The trained musician should see that these distances roughly correspond to perceived closeness. For instance, the second column is related to a particular inversion of the C minor chord. We see that the closest chord is the second inversion of the same chord, as described in Levine (1990). Hence,

<sup>3</sup>All notes with the same note name (e.g.  $C\#$ ) are said to be part of the same *pitch class*.

Table 1. Euclidean distances between the chord in the first row and other chords when chord representation is given by Equation (2).

$c1a2e3g3$	0.000	$c1d\#2a\#2d3$	0.000
$c1a2c3e3$	1.230	$c1a\#2d\#3g3$	1.814
$c1a2d3g3$	1.436	$c1e2a\#2d\#3$	2.725
$c1a1d2g2$	2.259	$c1a\#2e3g\#3$	3.442
$c1a\#2e3a3$	2.491	$c1e2a\#2d3$	3.691
$a0c3g3b3$	2.920	$a\#0d\#2g\#2c3$	3.923
$c1e2b2d3$	3.162	$a\#0d2g\#2c3$	4.155
$c1g2c3e3$	3.398	$g\#1g2c3d\#3$	4.363
$a0g\#2c3e3$	3.643	$c1e2a\#2c\#3$	4.612
$c1f2c3e3$	3.914	$a\#1g\#2d3g3$	4.820
$c1d\#2a\#2d3$	4.295	$f1a2d\#3g3$	5.030
$e1e2g2c3$	4.548	$d1f\#2c3f3$	5.267
$g1a\#2f3a3$	4.758	$a0c3g3b3$	5.473
$e0g2d3f\#3$	4.969	$g1f2a\#2c\#3$	5.698
$f\#0e2a2c3$	5.181	$b0d2a2c3$	5.902
$g\#0g2c3d\#3$	5.393	$e1d3g3b3$	6.103
$f\#1d\#2a2c3$	5.601	$f\#1e2a\#2d\#3$	6.329
$g0f2b2d\#3$	5.818	$d\#1c\#2f\#2a\#2$	6.530
$g1f2a\#2c\#3$	6.035	$g\#0b2f3g\#3$	6.746
$g1f2b2d\#3$	6.242	$b0a2d\#3g3$	6.947

we raise the note  $d\#3$  by one octave and replace the note  $d3$  by  $g3$  (separated by a perfect fourth.) These two notes are sharing some harmonics, leading to a close vectorial representation. This distance measure could have considerable interest in a broad range of computational generative models in music as well as for music composition.

### 3. Graphical Model

We now propose a graphical model that generates chord sequences using the input representation described in Section 2. The main assumption behind the proposed model is that conditional dependencies between chords in a typical chord progression are strongly tied to the metrical structure associated to it. Another important aspect of this model is that it is not restricted to local dependencies, like a simpler Hidden Markov Model (HMM) would be. This choice of structure reflects the fact that a chord progression is seen in this model as a two dimensional architecture. Every chord in a chord progression depends both on its position in the chord structure (global dependencies) and on the surrounding chords (local dependencies.) We show in Section 4 that considering both aspects leads to better generalization performance as well as better generated results than by only considering local

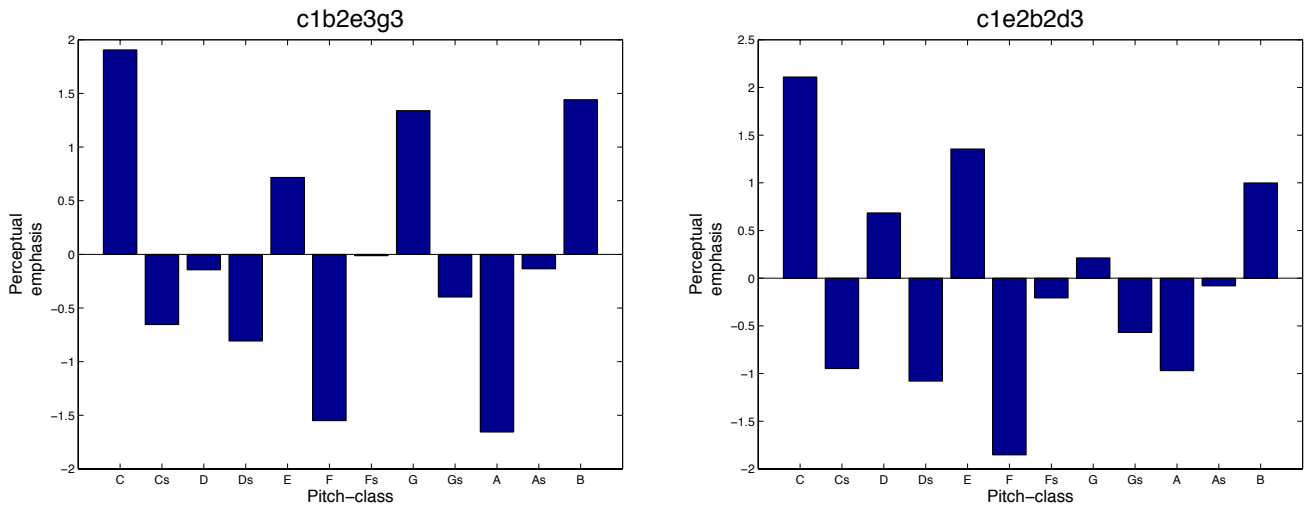


Figure 1. Normalized values given by Equation (2) for 2 standard 4-note jazz piano voicings of the C major chord. We see that perceptual emphasis is higher for pitch-classes present in the chord. These two chord representations have similar values for pitch-classes that are not present in either chords, which makes their Euclidean distance small.

dependencies. Figure 2 shows a graphical model that can be used as a generative model for chord progressions in this fashion.

Discrete nodes in Levels 1 and 2 are unobserved. The purpose of the nodes in Level 1 is to capture global chord dependencies related to the meter. Nodes in Level 2 are modeling local chord dependencies conditionally to the global dependencies captured in Level 1. For instance, the fact that the algorithm is accurately generating proper endings is constrained by the upper tree structure. The hierarchical structure imposed by vertical links is not sufficient for achieving smooth voice leadings (e.g. small distances between generated notes in two successive chords). This smoothness is modeled with the horizontal links in Level 2.

The bottom nodes of the model are continuous observations conditioned by discrete hidden variables. Hence, a mixture of Gaussians can be used to model each observation given by the distributed representation described in Section 2. Suppose a Gaussian node  $G$  has a discrete parent  $D$ , then the conditional density  $p(G|D)$  is given by

$$p(G|D = i) \sim \mathcal{N}(\mu_i, \sigma_i) \quad (3)$$

where  $\mathcal{N}(\mu, \sigma)$  is a  $k$ -dimensional Gaussian distribution with mean  $\mu \in \mathbb{R}^k$  and diagonal covariance matrix  $\Sigma \in \mathbb{R}^k \times \mathbb{R}^k$  determined by its diagonal elements  $\sigma \in \mathbb{R}^k$ .

The inference in the model can be done with the EM algorithm (Dempster et al., 1977) and the marginalization in the graphical model (during the expectation

step) can be achieved using the Junction Tree Algorithm (Lauritzen, 1996). Each  $\mu_i$  and  $\sigma_i$  in Equation (3) are learned using the EM algorithm. Exact marginalization techniques are tractable in this model given its limited complexity.

Many variations of this particular model can be possible, some of which are compared in Section 4. For instance, conditional probability tables can be tied in various ways. Also, more horizontal links in the model can be added to reinforce the dependencies between higher-level hidden variables.

It should be pointed out that in this paper we only consider musical productions with fixed length. Fortunately, the current model could be easily extended to variable length musical production by adding conditional dependencies arrows between many normalized subtrees.

## 4. Experiments

52 jazz standards excerpts from Sher (1988) were interpreted and recorded by the first author in MIDI format on a Yamaha Disklavier piano. See <http://cf.geocities.com/chordprogs> for a listing. Standard 4-note jazz piano voicings as described in Levine (1990) were used to convert the chord symbols into musical notes. Thus, the model is considering chord progressions as they might be expressed by a trained jazz musician in a realistic musical context. We chose to record actual voiced chords rather than symbolic chord names (e.g. Em7) because the sym-




Figure 2. A probabilistic graphical model for chord progressions. White nodes correspond to discrete hidden variables while gray nodes correspond to observed multivariate mixture-of-Gaussians nodes. Nodes in Level 1 directly model the contextual dependencies related to the meter. Nodes in Level 2 combine this information with local dependencies in order to model smooth chord progressions. Finally, continuous nodes in Level 3 are observing chords embedded in the continuous space defined by Equation (2). Numbers in Level 1 nodes indicate a particular form of parameter sharing that has proven to be useful for generalization (see Section 4.1).

bolic names are ineffective at capturing the specific voicings made by a trained jazz musician.

Every jazz standard excerpt was 8 bars long, with a 4 beat meter, and with one chord change every 2 beats (yielding observed sequences of length 16.) Longer chords were repeated multiple times (e.g. a 6-beat chord is represented as 3 distinct 2-beat observations.) This simplification has a limited impact on the quality of the model since generating a chord progression is simply a first (but very important) step toward generating complete polyphonic music, where modeling actual event lengths would be more crucial. The jazz standards were carefully chosen to exhibit a 16 bar global structure. We used the last 8 bars of each standards to train the model. Since every standard ends with a cadenza (i.e. a musical ending), the chosen excerpts exhibits strong regularities.

#### 4.1. Generalization

The chosen discrete chord sequences were converted into sequences of 12-dimensional continuous vectors as described in Section 2. Frequencies ranging from 20Hz to 20kHz (MIDI notes ranging from the lowest note in the corpus to note number 135) were considered in order to build the representation given by Equation (1). A value of  $\rho$  of 0.96 was arbitrarily chosen for the experiments. It should be pointed out that since the generative models have been trained in an unsupervised setting, it is irrelevant to compare different chord representations in terms of likelihood. However, it is pos-

Table 2. Mean negative out-of-sample log-likelihoods of sub-sequences of length 4 on positions 1, 5, 9 and 13. These results are computed using double cross-validation in order to optimize the number of possible values for hidden variables. The numbers in parentheses indicate which levels of the tree are tied as described in Section 4.1. We see that some combinations of parameter tying in the trees performs better than the standard HMM.

Model (tying)	Negative log-likelihood
Tree (2, 3)	93.8910
Tree (1, 3)	94.0037
Tree (1, 2, 3)	94.9309
Tree (3)	98.2446
HMM	98.2611

sible to measure how well a given architecture is modeling conditional dependencies between sub-sequences of chords. In order to do so, mean negative out-of-sample likelihoods of sub-sequences of length 4 on positions 1, 5, 9 and 13 have been computed. Double cross-validation has been used to optimize the number of possible values of hidden variables for various architectures. Results are given in Table 2.

Different forms of parameter tying for the tree model shown in Figure 2 have been tested. All nodes in Level 3 share the same parameters for all tested models. Hence, we use only one 12-dimensional mixture of Gaussians (as in Equation (3)) independently of time,

which serves to constrain the capacity of the model. Moreover, a diagonal covariance matrix  $\Sigma$  has been used, thus reducing the number of free parameters to 24 in Level 3 (12 for  $\mu$  and 12 for  $\Sigma$ ). Hidden variables in Level 1 and 2 can be tied or not. Tying for Level 1 is done as illustrated in Figure 2 by the numbers inside the nodes.

The fact that the contextual out-of-sample likelihoods presented in Table 2 are better for the different trees than for the HMM indicates that time-dependent regularities are present in the data. Sharing parameters in Levels 1 or 2 of the tree increases the out-of-sample likelihood. This indicates that regularities are repeated over time in the signal. Further investigations would be necessary in order to assess to what extent chord structures are hierarchically related to the meter.

The relatively high values obtained in terms of conditional out-of-sample negative log-likelihood indicates that the number of sequences may not be sufficient to efficiently represent the variability of the data. Unfortunately, reliable chord progressions data is difficult to generate. However, the authors plan to extend the chord database significantly in the short term.

## 4.2. Generation

One can sample the proposed model in order to generate novel chord progressions. Fortunately, Euclidean distances are relevant in the observation space created in Section 2. Thus, a simple approach to generate chord progressions is to take the nearest neighbors (in the training set) of the values obtained by sampling the observation nodes.

Chord progressions generated by the models presented in this paper are available at <http://cf.geocities.com/chordprogs>. For instance, Figure 3 shows a chord progression that has been generated by the graphical model shown in Figure 2. This chord progression has all the characteristics of a standard jazz chord progression. For instance, the trained musician can observe that the last 8 bars of the sequence is a II-V-I<sup>4</sup> chord progression (Levine, 1990), which is very common.

Figure 4 shows a chord progression generated by the HMM model. While the chords do follow one another in a smooth fashion, there is no global relation between chords. For instance, one can see that the lowest note

---

<sup>4</sup>The lowest notes are d, g and c corresponding to the second, fifth and first notes of the major scale.

of the last chord is not a c, which was the case for all the chord sequences in the training set. The fundamental qualitative difference between both methods should be obvious even for the non-musician when listening to the generated chord sequences.

## 5. Conclusion

We have shown empirically that chord progressions exhibit global dependencies that can be better captured with a tree structure related to the meter than with a simple dynamical HMM that concentrates on local dependencies. The importance of contextual information for modeling chord progressions is even more apparent when one compares sequences of chords sampled from both models. The time-dependent hidden variables enable the tree structure to generate coherent chord progressions both locally and globally.

However, the low difference in terms of conditional out-of-sample likelihood between the tree model and the HMM, and the relatively low optimal capacity for generalization, are good indications that increasing the number of sequences in the dataset would probably be necessary in further developments of probabilistic models for chord progressions. Also, a better evaluation of such models could be achieved by exploring supervised tasks such as automatic transcription, music information retrieval, musical genre recognition, and music analysis applications, just to name a few.

Chord progressions are regular and simple structures that condition dramatically the actual choice of notes in polyphonic tonal music. Hence, we argue that chord models are crucial in the design of efficient algorithms that deal with such music data. Moreover, generating interesting chord progressions may be one of the most important aspects in generating realistic polyphonic music. Our model constitutes a first step in that direction.

## References

- Allan, M., & Williams, C. K. I. (2005). Harmonising chorales by probabilistic inference. *Advances in Neural Information Processing Systems*.
- Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5, 157–166.
- Cemgil, A. T. (2004). *Bayesian music transcription*.




Figure 3. A chord progression generated by the proposed model. This chord progression is very similar to a standard jazz chord progression.




Figure 4. A chord progression generated by the HMM model. While the individual chord transitions are smooth and likely, there is no global chord structure.

Doctoral dissertation, Radboud University of Nijmegen.

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society*, 39, 1–38.

Eck, D., & Schmidhuber, J. (2002). Finding temporal structure in music: Blues improvisation with LSTM recurrent networks. *Neural Networks for Signal Processing XII, Proc. 2002 IEEE Workshop* (pp. 747–756). New York: IEEE.

Handel, S. (1993). *Listening: An introduction to the perception of auditory events*. Cambridge, Mass.: MIT Press.

<http://cf.geocities.com/chordprogs> (2005). Temporary anonymous URL to put audio samples generated by the proposed algorithms. A proper academic website will be established if the paper is accepted.

Kuusi, T. (2001). *Set-class and chord: Examining connection between theoretical resemblance and per-*

*ceived closeness*. No. 12 in *Studia Musica*. Sibelius Academy.

Lauritzen, S. L. (1996). *Graphical Models*. Oxford University Press.

Lavrenko, V., & Pickens, J. (2003). Polyphonic music modeling with random fields. *Proceedings of ACM Multimedia*. Berkeley, CA.

Levine, M. (1990). *The Jazz Piano Book*. Sher Music Co./Advance Music.

Moore, B. (1982). *An Introduction to the Psychology of Hearing*. Academic Press.

Raphael, C., & Stoddard, J. (2003). Harmonic analysis with probabilistic graphical models. *Proceedings of ISMIR 2003*.

Sher, C. (Ed.). (1988). *The New Real Book*, vol. 1. Sher Music Co.

Valimaki, V., Huopaniemi, J., Karjalainen, & Janosy, Z. (1996). Physical modeling of plucked string instruments with application to real-time sound synthesis. *J. Audio Eng. Society*, 44, 331–353.

Vassilakis, P. (1999). Chords as spectra, harmony as timbre. *J. Acoust. Soc. Am.*, *106*, 2286. (paper presented at the 138th meeting of the Acoustical Society of America).