



II V I and all that Jazz

Description

With more time spent teaching in front of a computer I've learnt more about developments in automated speech recognition. Software that turns speech into text (as transcriptions and closed captions) is a major accomplishment. Practically it's artificial intelligence (AI). Researchers attempt the same with music: turning an audio music track into musical notation: notes on a page, chord charts, MIDI notation, etc.

A book by Meinard Müller provides a comprehensive account of the techniques and challenges of turning audio into notation. How can an automated system identify notes played at the same time, and thereby identify chords and sequences of chords? That's one of many challenges for such systems.

Researchers have developed programs for breaking audio signals into frequency bands using Discrete Fourier Transform methods (related to [DCT](#) methods I've described in previous posts). However, further steps are required to identify and name the chords you might hear on an audio track.

Audio signals are ambiguous. Not least, single notes carry overtones depending on the instrument that is playing it, and different chords can generate similar spectral distributions depending on the varied intensities of the notes in the chord.

Müller calls the spectral output from a particular chord, once it has been filtered, a "chroma-based audio feature". It's a coloured spectrogram showing bands of colour differentiating prominent frequencies (i.e. notes). It's easy enough to generate the most likely spectral output from any chord played on any instrument simply by playing (or simulating) the instrument. It's more difficult to reverse engineer this process i.e. identify the chord from the audio.

Exploiting musical patterns

The solution lies in methods for predicting what is the chord most likely to appear at any moment in an audio track of a piece of music. For all its inventiveness, tonal music from pop to classical, exhibits chord sequences (progressions) that are repeated within any composition, but also from one piece of music to the next.


To illustrate, a music group calling themselves [The Axis of Awesome](#) published a beguiling YouTube video with the tagline “Ever wonder why all those pop songs sound kinda the same? Well, it’s pretty simple; They all use the same 4 Chords!” The progression they refer to is I, V, VI and IV, i.e. triads based on the first, fifth, sixth and fourth notes of a major scale. In the key of C major that’s the chords C, G, Am and F in that order and repeated.

They amalgamate *Let It Be* (Beatles), *Take me Home*, *Country Road* (Denver), *No Woman, No Cry* (Marley), *I Come from a Land Down Under* (Hay, Strykert), *You’re So Beautiful* (Bostwick, Smollett, Washington), *Time to say Goodbye* (Sartori, Peterson, Quarantotto), *Auld Lang Syne*, etc. The songs are dramatic, if not [melancholic](#).

Jazz chord progressions

As lock-down therapy, I re-familiarised myself with guitar chord progressions. Instructional videos kept reminding me that most jazz improvisation is founded on the II, V, I chord progression, e.g. Dm, G, C. Another feature of jazz is that it works less with musical triads (3 note chords) and mostly with 4 note chords. The extra note creates a disharmony and tension. So the Dm, G, C chordal progression becomes Dm7, G7 and Cmaj7.

An automated chord recognition algorithm can exploit such predictable sequences to improve its chances of identifying the chords in the audio. That’s where the [Hidden Markov Method](#) (HMM) comes into play. Here’s a network diagram showing a Markov model with various chord sequences.



The II, V, I progression is shown as the grey circles and their arrows. The green arrows connect the I, V, VI, IV progression (*Let it Be*, etc). The heavy arrows connect a more complicated but common classical and jazz sequence: II, V, I, IV, VII, II, VI. (The jazz standard *Autumn Leaves* (Kosma, Mercer, Prevert) uses that.) The purple chords belong to the related minor scale. The *transition table* shows the probability that each of the chords on the first column is followed by the chord indicated in the top row. The numbers are effectively probability labels on the arrows in the network.

Hidden and observed

I derived these numbers by guesswork, but something similar can be derived by inspecting the chord progressions evident in a corpus of musical pieces, e.g. (jazz standards from *The Real Book*). Some of the authors of the references below draw on databases of hundreds of instances of chord sequences to derive such probabilities.


Note also that it's unlikely there is zero probability that one chord is followed by another chord. I've omitted small probabilities from this diagram for simplicity. One could imagine a typology of such relationships depending on musical genre, style, performer, etc. Different styles exhibit different chord sets (nodes), and different probability values. Key changes add extra complexity.

Such diagrams are Markov models, as described in a [previous post](#). They are "hidden" in the sense that we don't know their actual instantiation in any particular musical piece when all we have is the audio.

Markov models can also generate sequences. If the diagram above were part of a system for generating music, then an automated chord player would jump from node to node randomly, and in a way that is biased according to the probabilities along the links.

The "unhidden" aspect of this approach to chord detection is the *observation* part of the model: the probabilities that any known chord will deliver a particular spectral distribution. Calculations with these relationships and probabilities increases the chances that a system for converting music audio into musical notation, or at least a chord chart, will hit on the right chords, accurately. M ller explains the algorithms for this.

Here are the chords in the key of C major: classical/folk versions on the left. The jazz version is on the right.



References

- Berget, Gunhild Elisabeth. 2017. *Using Hidden Markov Models for Musical Chord Prediction (MSc thesis)*. Norwegian University of Science and Technology Department of Mathematical Sciences, Trondheim
- Kiefer, Peter, and Manda Riehl. 2016. Markov Chains of Chord Progressions. *Ball State Undergraduate Mathematics Exchange*, (10) 1, 16-21.
- M  ller, Meinard. 2015. *Fundamentals of Music Processing: Audio, Analysis, Algorithms, Applications*. Heidelberg: Springer
- Various. 2014. *The Real Book Volume I (6th Edition)*. Milwaukee: Hal Leonard



Category

1. Music

Tags

1. AI
2. chords
3. guitar
4. Markov
5. music

Date Created

October 10, 2020

Author

rcoyne99