APP MTH 3001 Applied Probability
Group project: Turning Charlie Parker into Ornette Coleman
Due: 3pm Friday 18 May

Preamble

This project is about algorithmic music composition, specifically, about creating original jazz improvisations from existing pieces of music.

To do this, we'll use two common file formats used by musicians to store musical information:

- *MusicXML*: an open file format for storing melodies and chord information (more on that soon);
- *MIDI*: for storing information on which notes to play at which times, for playing in GarageBand or similar software.

You'll be given a transcription of a melody + improvisation by the jazz saxophonist Charlie "Bird" Parker (one of the originators of "bebop", and therefore, modern jazz music) in MusicXML format, and asked to develop an algorithm (based on Markov Chains) to generate new improvisations. We'll explore this in a few different ways, but first, a crash course on music and jazz theory.

Crash course in music

At its simplest, a musical melody consists of a sequence of notes (12 semi-tones per octave) and the time at which that note should be played. MIDI files denote each note by an integer from 0 to 127, while MusicXML denotes each note by the octave + number of semitones within this octave (e.g., note 37 in MIDI format would be stored as 3 octaves + 2 semitones in XML). We shall also ignore rhythm completely and consider only the sequence of notes played, not the timing. However, in your project you are welcome to improve your models to take into account some of these features.

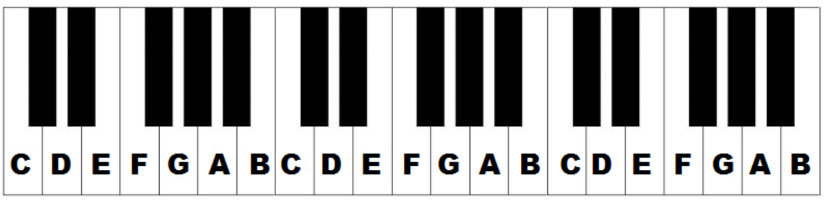

Beyond melody, we have *chords*: groups of notes played at the same time as the melody to provide harmony. Different chords give different tonalities to the music (e.g., "major" chords sound happy, while "minor" chords sound sad), and an ordering of chords in time is called a *chord progression*, which gives temporal structure to a piece of music. There are many common types of chord progressions[1] (you'll see some of them in this project), but the most important point is: **not all 12 unique notes should be played over each chord! Some notes sound very bad over different chords**.

In this project we'll explore the interaction between melodies and chords.

Crash course in jazz

The above covers from the beginning of time through to about the 1940s, so now let's describe the structure of *jazz music*. Most jazz performances follow this standard structure, played over a repeating chord progression:

1. An exposition of the melody (the "head"), usually played over the full sequence of chord changes in the song, with all instruments playing.
2. Improvisation(s) over the same sequence of chord changes, generally with different instruments each taking turns to play a "solo". These immprovisations are often variations on the head.
3. A recapitulation of the head, as in part 1, to bookend the performance[2].

Common chord progressions are a 12-bar "blues", or a 32-bar "rhythm changes" (this has its own internal structure, similar to modern pop music: melody + melody + bridge + melody).

Over the chords to Gershwin's "I Got Rhythm".

The innovation of Charlie "Bird" Parker and other modern jazz musicians was not any of these "rules" or chord progressions, which had been around since before the 20th century, but the types of notes ("scales") played over the chords. Parker and contemporaries played many blues' and rhythm changes', but the types of notes played were new and fresh-sounding. However, the fundamental principles remain: certain notes don't sound good over certain chords.

As Parker's recordings are now legendary in jazz history, many of his compositions + solos have been transcribed and immortalised in the so-called "Omnibook". These transcriptions will form the basis of our experimentations in algorithmic music composition.

The project

The aim of the project is to generate new improvisations in the style of Charlie Parker. We'll explore a few different ideas to do this. (These

[1]For example, you might be aware that most pop songs are made from just four chords.

[2]Bill Evans explains all of this much better than I do!

are not even remotely exhaustive!) Each group will be provided with their own piece from the Omnibook to analyse. You will be provided with an XML file of the particular transcription you should be considering. To import into Matlab and obtain the sequence of notes played, you will need to perform some "data wrangling" on this file. I used two set of scripts:

- Eric Nichols' MusicXML parsing scripts from the University of Indiana (`http://music.informatics.indiana.edu/code/musicxml/`), to convert XML into a Matlab `.mat` file;
- Toiviainen & Eerola's Matlab MIDI toolbox project (`https://github.com/MIDItoolbox/1.1`), to export music generated in Matlab to MIDI.

Yes, you will need to learn about Matlab's `mex` function. Mac users: you may learn something about how OSX vs. Windows describes paths to files!

You should answer the following questions in the form of a project report that is to be well-written, word processed and spell checked with any supporting material being appropriately presented. Treat this project as if I am your boss and I have asked your group to answer these questions and write me an appropriate report that I can use in dealing with a jazz musician client. My meeting with the client has been scheduled for 9am on Monday 21 May. Hence I need the report by 3pm on Friday 18 May. As I will have very little time to check your results, I need your report to be written in such a way as to demonstrate why your results are correct.

If you want to use LaTeX, I'll make a report template available, and can recommend Overleaf for collaboration.

Approach 1

Assume that each new note is chosen randomly, depending only on the current note, and independently of all previous notes.

a) Define your own state space $\mathcal{S}$. This might be all the notes available, thus $\mathcal{S} = \{0, 1, 2, \ldots, 127\}$, or it might be only the notes played in your piece of music, or something else entirely. For simplicity, in what follows, I shall refer to the elements of $\mathcal{S}$ as notes, and the sequence of sounds played as the sequence of notes, although these may not be strictly true.

b) Argue why this stochastic process is Markovian.

c) Use the piece of music you were given to determine the proportion of times that the MIDI key j immediately follows the MIDI key i, for all $i, j \in \mathcal{S}$. What did you do with the last note in your piece of music? Think of at least two ways you could treat this note, and discuss the fundamental consequences of your choices.

d) Use this to construct a $\mathbb{P}$-matrix representation for the evolution of your piece of music, including all states $i \in \mathcal{S}$.

   i) Are there any zero columns, and if so, explain why and say how many? What do the associated rows look like? Do you have any choice in this?

ii) Ignoring any states associated with zero columns, classify the character of all remaining states.

e) Let $\mathcal{S}_+$ denote the subset of $\mathcal{S}$ for which the proportion is strictly positive. What are the probabilities $P(X_1 = K \big| X_0 = i)$ for $i \in \mathcal{S}_+$?

f) What are the probabilities $P(X_{10} = K \big| X_0 = i)$ for $i \in \mathcal{S}_+$?

g) Write down the equations that determine the expected time until the note K is first played, given that you start your music with note $i \in \mathcal{S}$.

h) Solve these equations for all $i \in \mathcal{S}$.

i) How do you know whether your values are correct?

j) Create a new solo for your piece of music, by simulating a sequence of notes from your stochastic process. Use the MIDI toolbox to play your solo *over the chord progression for your piece*, and convert it into an audio file so that you can listen to the piece of music.
Comment on your new solo.

Approach 2

Now let's explicitly incorporate the notion that certain notes "go with" certain chords, by considering the evolution of note-chord pairs. Assume that each new note-chord pair is chosen randomly, depending only on the current note **as well as** the current chord being played.

a) Define an appropriate state space $\mathcal{T}$ for this stochastic process, such that it is Markovian. Give the details of the new state space, $\mathcal{T}$, and explain why this stochastic process is Markovian on $\mathcal{T}$.

b) Use this to construct a $\mathbb{P}$-matrix representation for the evolution of your piece of music.

c) Write down the equations that determine the expected time until the note-chord combination K is first played, given that you start your music with the note-chord ordered pair $(i, j), (i, j) \in \mathcal{T}$.

d) Solve these equations for all ordered pairs $(i, j) \in \mathcal{T}$.

e) How do you know whether your values are likely to be correct?

f) Simulate a sequence of notes and chords from your stochastic process and convert it into an audio file so that you can listen to it. Comment on this new solo. How is it different to the original piece, and to Approach 1?

Approach 3

The final extension is to explicitly consider: notes, chords, and the chord progression for your transcription jointly. Assume that each new note-chord pair is chosen randomly, depending only on the current note as well as the current chord being played, **and** the chord progression follows that of the original piece.

a) Argue for why this process is now a time-inhomogeneous Markov Chain.

b) Construct a $\mathbb{P}$-matrix representation for the evolution of your piece of music. Make sure it is explicit in your report how you have stored $\mathbb{P}$.

or $\mathbb{P}$-matrices.

c) Simulate a new solo over the original chord progression from your stochastic process. Make sure you detail any assumptions or choices you had to make that are different to the previous Approaches in order to produce your piece.

d) Comment on this solo and how it sounds different to those from Approach 1 & Approach 2.

Final notes

- While this document is set out like a standard (long) assignment, you'll find as you get into it that the instructions are not totally prescriptive about everything you need to do, and that you'll have to make various *choices* and assumptions along the way. You should make sure that your final report clearly records each of these choices (along with some explanation for why you made the choices you did). Think broadly about *reproducible research*: I should be able to reproduce essentially what you did, by reading your report.
- Keep in mind also that I'm going to be pitching your work to a client, so you should make sure your final "product" sounds good! I'll leave the interpretation of this up to you.
- Week 10 will come around sooner than you think, so get started early. I will prompt you throughout the semester by asking you questions related to this project on your assignments, which will give you an idea of where you should be up to.
- No need to be worried if you're not a musician – it's designed to be doable even if you know nothing about music. That said, I've written this having some (very limited) knowledge of music/jazz theory, so if **anything** is unclear, please don't hesitate to come and ask. I'm looking forward to chatting/explaining/mentoring about this project throughout the semester.
- Finally, this is intended as a somewhat open-ended project, so don't be afraid to **play**, explore, and have fun!