VISUALISING CHORD PROGRESSIONS IN MUSIC COLLECTIONS: A BIG DATA APPROACH

Alexander Kachkaev[1] , Daniel Wolff[1] , Mathieu Barthet[2], Dan Tidhar[1], Mark Plumbley[2], Jason Dykes[1], Tillman Weyde[1]

[1] *City University London, Department of Computer Science*

[2] *Queen Mary University of London, Centre for Digital Music*

Correspondence should be addressed to: alexander.kachkaev@city.ac.uk

Abstract: **In the Digital Music Lab project we work on the automatic analysis of large audio databases that results in rich annotations for large corpora of music. The musicological interpretation of this data from thousands of pieces is a challenging task that can benefit greatly from specifically designed interactive visualisation. Most existing big music data visualisation focuses on cultural attributes, mood, or listener behaviour.**

In this ongoing work we explore chord sequence patterns extracted by sequential pattern mining of more than one million tracks from the *I Like Music* commercial music collection. We present here several new visual representations that summarise chord patterns according to chord types, chroma, pattern structure and support, enabling musicologists to develop and answer questions about chord patterns in music collections.

Our visualisations represent root movement and chord qualities mostly in a geometrical way and use colour to represent pattern support. We use two individually configurable views in parallel to encourage comparisons, either between different representations of one corpus, highlighting complimentary musical aspects, or between different datasets, here representing different genres. We adapt several visualisation techniques to chord pattern sets using some novel layouts to support musicologists with their exploration and interpretation of the corpora.

We found that differences between chord patterns of different genres, e.g. Rock & Roll vs. Jazz, are visible and can be used to generate hypotheses for the study of individual pieces, further statistical investigations or new data processing and visualisation. Our designs will be adapted as user needs are established through ongoing work. Means of aggregating, focusing and filtering by selected characteristics (such as key, melodic patterns etc.) will be added as we develop our design for the visualisation of chord patterns in close collaboration with musicologists.

The visualisations are available as a web application at `http://dml.city.ac.uk/csvd/`

1. INTRODUCTION

Our interest in the Digital Music Lab[1] project lies in the analysis of large corpora of music, exposing their specificities, for example across different genres, as well as determining the commonalities of the pieces within a corpus. To derive rich annotations at the collection-level, we use automatic analysis methods from the field of Music Information Retrieval in conjunction with Big Data methods that enable fast and efficient analysis. The analysis of the resulting data is challenging. Graphical representations of music and its structures can provide valuable tools for musicologists and music research.

The musicological interpretation of the data resulting from thousands of pieces requires an adequate representation of this data. Various techniques exist for the analysis of harmony and chord sequences, such as chord charts, annotations of chord functions, but most conventional graphical representations are restricted to the scope of a single piece and do not trivially extend to collections of pieces.

Although recently more methods of collection-level visualisation have been introduced, the available analysis focuses on cultural attributes, mood, or listener behaviour whereas we are focusing on attributes closer to the notation of the music itself (e.g. chords, keys, transcriptions, etc.).

[1] `http://dml.city.ac.uk/`

In this paper we present visualisations that present frequent chord sequence patterns obtained for different genres of popular music. Apart from determining the most frequent chord sequences, the different visualisations can highlight various aspects of the analysis and support for comparison across different genres.

2. THE DATASETS

In our ongoing work we explore chord sequence patterns, i.e. frequently occurring parts of the chord progressions. The pattern data contain the chord sequence as it has been identified in the music as well as the *support* value, which is the percentage of pieces in which a pattern appears.

The sequences have been extracted using sequential pattern mining (the CM-SPADE algorithm [1]) on more than one million tracks of the "I Like Music" commercial music collection [2]. See [10] for details of the extraction process. In this paper we are looking specifically at genre-defined subsets, which are listed in Table 1.

Table 1: Genre subsets with numbers of tracks.

Genre	*Tracks*
Jazz	35991
Rock & Roll	36654
Classical	21446
Blues	31618
Reggae	13421
Folk	45194
Rap	24880

The most patterns are extracted from the genre-defined subsets. Table 2 shows the sequences of two or three chords with highest support in the *Blues* subset.

Table 2: Common chord sequences for the Blues collection.

Chord 1	Chord 2	Chord 3	...	Support
G	C	–	–	13.134923
C	G	–	–	12.973623
C	F	–	–	11.771776
C	C7	–	–	11.768613
F	C	–	–	11.616801
...				
G	C7	C	–	6.923272
D	Gm	–	–	6.920109
C7	C	G7	–	6.920109
...				

3. THE VISUALISATION

We are designing new visual methods that summarise chord pattern sets with regard to chord types, roots, structure and support, enabling musicologists to develop and answer questions about the

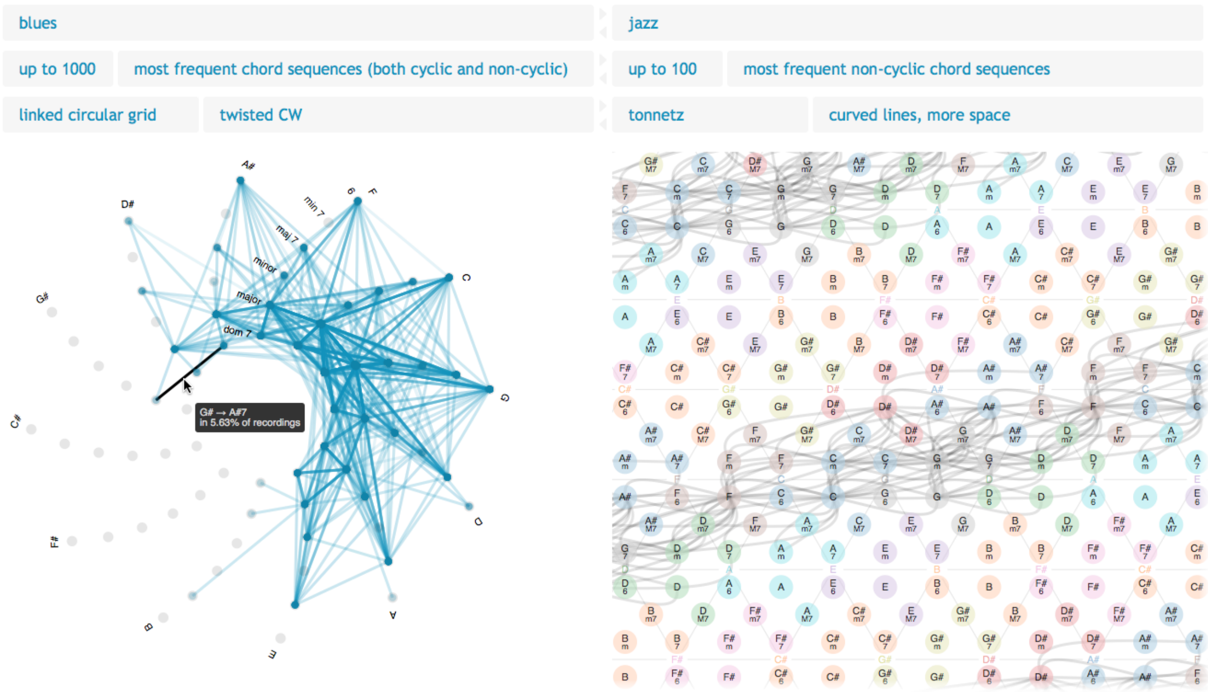

Figure 1: Interface of the chord progressions visualising tool. Available at `http://dml.city.ac.uk/csvd/`.

use of chord patterns in different genres. All data visualisations we have developed are accessible through a single user interface, the main tasks of which are the following:

- to let the user observe the distribution of patterns in the datasets;
- to support experimentation with various data layouts;
- to support the visual comparison of different datasets and visualisations.

3.1. *Interface Structure*

The interface is implemented as a web page, a snapshot of which is shown in Figure 1. It consists of two independent panels with controls to choose music genre, number of chord sequence patterns to show, types of sequences (non-cyclic, cyclic or both), the visualisation technique to use and its configuration.

This layout encourages the comparison between different visualisations of a corpus, which can highlight complementary musical aspects, as well as the comparison between datasets, in this case the genre-defined subsets.

3.2. *Visualisation techniques*

We have designed and apply five different visualisation techniques, which show chord sequences from different perspectives (see Figure 2).

Tonnetz-based diagram is a layout inspired by [6]. The chords here are arranged according to the proximity of the notes they consist of, which in turn form a lattice, originally invented by Leonhard Euler [4]. Chord sequences are displayed as linear connections between the nodes.

Linked grid is a tabular structure, where root notes are listed as columns, and chord types are put into rows. A chord sequence is visualised as a link between corresponding cells, and its frequency is mapped to opacity.

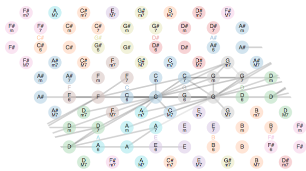

(a) Tonnetz-based diagram

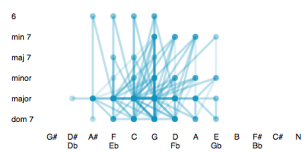

(b) Linked grid

(c) Circular linked grid

(d) Origin-destination matrix

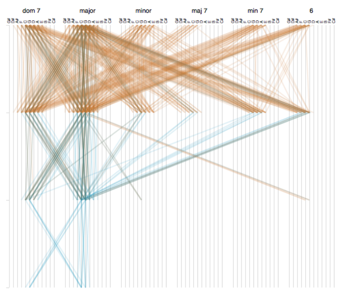

(e) Parallel coordinate plot

Figure 2: Designed visualisation layouts

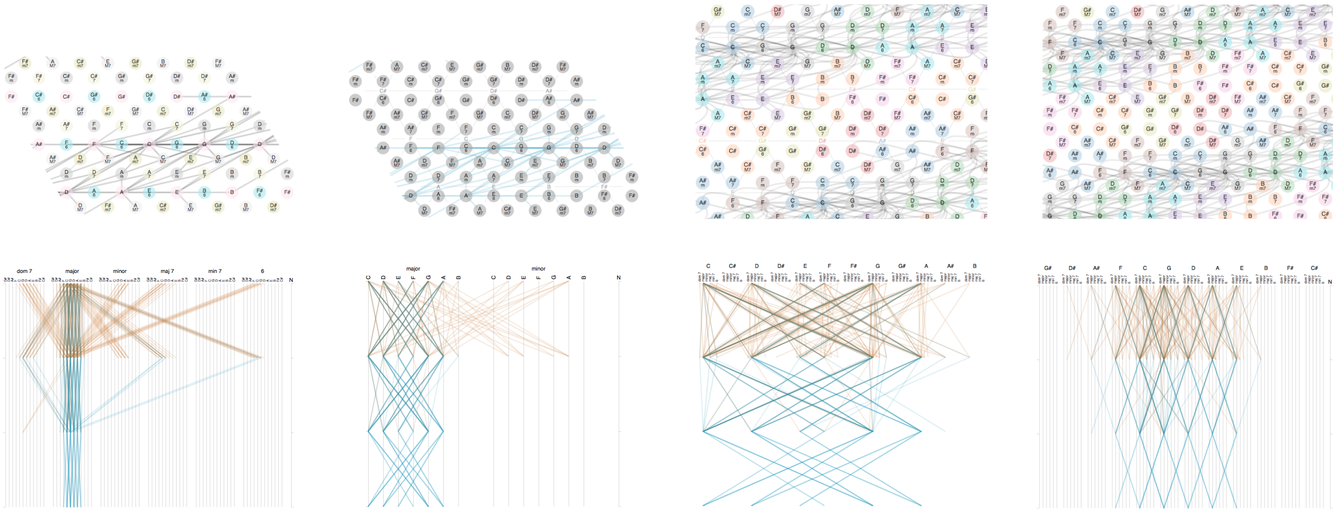

Figure 3: Examples of layout adjustment for Tonnetz-based diagram *(top)* and Parallel coordinate plot *(bottom)*. All views show two hundred most frequent chord progressions in Rock & Roll.

Linked circular grid is similar to *linked grid*, but with the polar coordinate system instead of the Cartesian one. This layout reduces overlaps of some links and lowers the probability of a ‘breaching effect’, which may occur when a chord sequence connects the first and the last column of the grid.

Origin-destination matrix is a visualisation layout, which has been commonly utilised for transport-related data since 1955 [3]. It can only be applied for progressions with two chords. The first chord (origin) corresponds to a row, and the following one (destination) is mapped to a column. A cell that is located at their intersection is coloured according to the support of the pattern.

Parallel coordinate plot is a common way of visualising multivariate data [5]. Chord names are placed along a horizontal axis and their positions in the progressions on the vertical axis, so that it is possible to display all sequences including ones with more than two chords.

3.3. Adjustment of layouts

The visualisation layouts are customisable to facilitate exploration of the datasets. For every visualisation technique it is possible to choose between four to six configurations, changing the order of chords in grids, meaning of colour, curvature and angle of lines, etc., as shown in examples in Figure 3. Although use of presets instead of controls for changing individual layout parameters does not provide full freedom to users, this approach significantly simplifies the interface and makes it easy to navigate between successful combinations of settings.

3.4. Interaction

Our interface currently provides two types of user interaction: (1) changing data or their representation in one of two panels and (2) getting contextual details on demand.

Switching between genres, types of chords, sample sizes, visualisation techniques and their configurations is done with drop-down menus on top of each panel, which can be controlled using mouse or keyboard and also serve as indicators of the current state. These controls are supplemented with pairs of shortcut buttons between the panels that allow states of the panels to be synchronised in at most three mouse clicks.

Contextual details on demand are brought to users in a form of a tool tip, which appears near a chord sequence when it is hovered over with a mouse (see Figure 1). The tooltip shows the names of the chords in the selected progression and the corresponding support.

Any change of the state of the interface is reflected in the URL. First, this makes it possible to navigate between recently opened views using the ‘back’ and ‘forward’ functions in the browser history, thus reducing the cost of navigation. Second, this allows users to share direct links to the views and better collaborate with colleagues [9].

3.5. Technologies

Our interface uses open web technologies and can be used in all modern browsers on desktop, laptop and tablet devices. Most of the program code is executed on the client side and is written in JavaScript. Rendering of visualisations is done with help of Vega[2], a visualisation grammar that works on top of D3[3]. All generated graphics are in vector format, so they can be easily zoomed for detailed exploration or exported for high-resolution printing. Handling of user actions and other events is organised using Marionette.js[4], a composite application framework that links Backbone.js[5], Underscore[6] and jQuery[7] – JavaScript libraries that simplify software development and improve browser compatibility.

The server-side components of the application is implemented with Symfony[8], a general-purpose PHP framework. Although the only task of the server in this particular case is to deliver all necessary resources to the front-end, Symfony plays an important role in code structuring and back-end optimisation. For example, it helps combine and compress style sheets and JavaScript files with Assetic[9], which significantly reduces the application launch time.

4. Application

We currently explore the benefits and limitations of our approach with preliminary user testing. So far, we found that the visualisations do show differences between chord patterns of different genres. These visible differences can be used to generate hypotheses for the study of individual pieces, further statistical investigations or new data processing and visualisation approaches. In this way, the proposed visualisations can, for instance, complement and support related methods for automatic genre classification (e.g. [7]). Our designs will adapt as user needs are established through ongoing work. Means of aggregating, focusing and filtering by selected characteristics (such as key, melodic patterns etc.) will be added as we develop our visualisation designs and expand the datasets in close collaboration with musicologists.

[2] `http://trifacta.github.io/vega/`
[3] `http://d3js.org/`
[4] `http://marionettejs.com/`
[5] `http://backbonejs.org/`
[6] `http://underscorejs.org/`
[7] `http://jquery.com/`
[8] `http://symfony.com/`
[9] `http://github.com/kriswallsmith/assetic`

(a) Blues vs. Jazz + linked circular grid

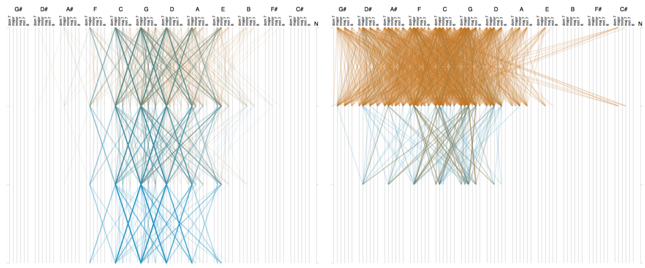

(b) Folk vs. Jazz + parallel coordinate plot

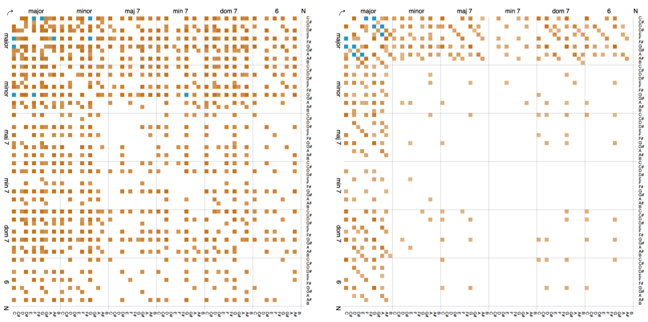

(c) Jazz vs. Rock & Roll + origin-destination matrix

Figure 4: Examples of revealed differences in datasets

4.1. Examples

We work on these visualisation techniques with musicologists, who find them useful to explore differences in harmonic structures between the genre-defined collections. This is illustrated in the following examples.

Figure 4a shows the comparison between the 100 most frequent Blues and Jazz chord sequence patterns using the circular grid view. The figures show that in the Blues collection chords tend to be further to the left of C in the circle of fifth. Blues shoes less use of patterns containing minor chords. Interestingly, this collection shows no patterns among the 100 most frequent in Jazz that contain *major 7* chords, while *major 7* chords do appear in Blues, contrary to common concepts of Jazz harmony (e.g. [8, pp. 51]).

Figure 4b shows the Folk and Jazz collections using the parallel coordinates view. A visually striking difference is that there is more diversity in the Jazz transitions and that frequent patterns tend to be short. On the other hand, in the Folk collection, there are many longer patterns (3 and 4 chords) with strong support , mostly cyclic (shown in blue) and moving in fifths (visible by the angles of the lines).

Figure 4c shows the Jazz and the Rock & Roll collections in an origin-destination matrix. This visualisation includes only 2-chord patterns and for Rock and Roll it shows that almost all patterns start or end on a major chord (all but the 1st row and column are very sparsely populated).

5. CONCLUSIONS

We have presented visualisations for chord sequence patterns extracted from large sets of music recordings. The presented approach allows the user to compare datasets and visualisations in various combinations, leading to interesting observations, such as the distribution of patterns containing major 7 chords between the Jazz and Blues subsets.

We propose this tool as a component of a larger system, the Digital Music Lab, which will include analysis of music corpora with regard to other musical parameters and metadata, such as temporal any geographic information. Our approach demonstrates that visualisation is a useful method to investigate datasets that are too large for direct inspection of all individual recordings.

5.1. Future Work

The next step for this particular application is to perform key detection, so that we can express and visualise harmonic progressions in terms of scale degrees or harmonic functions on large corpora, which is more closely related to music theory. More generally, we would like to extend the visualisation to other aspects of music such as rhythmic and melodic patters, as well as integrating the visualisation of individual music recordings.

ACKNOWLEDGEMENTS

This work was supported by the UK Arts and Humanities Research Council (Grant AH/L01016X/1)

REFERENCES

[1] Fournier-Viger, P., Gomariz, A., Campos, M., & Thomas, R. Fast Vertical Sequential Pattern Mining Using Co-occurrence Information. In *Proc. 18th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2014.

[2] Fazekas, G., Barthet, M. & Sandler, Mark B. The BBC Desktop Jukebox Music Recommendation System: A Large Scale Trial With Professional Users In *Proc. of the Int. Conf. on Multimedia and Expo (ICME)*, 2013.

[3] Voorhees, A. A general theory of traffic movement. In *Institute of Traffic Engineers Past Presidents' Award Paper*, New Haven 1955.

[4] Euler, L. Tentamen novae theoriae musicae ex certissismis harmoniae principiis dilucide expositae. Saint Petersburg Academy, 1739.

[5] Inselberg, A., & Dimsdale, B. Parallel coordinates: a tool for visualizing multi-dimensional geometry. In *Visualization'90., Proceedings of the First IEEE Conference on. IEEE*, 1990.

[6] Tymoczko, D. The Generalized Tonnetz. In *Journal of Music Theory*, 2012.

[7] Pérez-Sancho, C., Rizo, D., & Iñesta, J. M. Genre classification using chords and stochastic language models. In *Connection science*, 21(2-3), 145-159

[8] Nettles, B. Harmony 2. Berkeley College of Music, 1987.

[9] Walker, R., Slingsby, A., Dykes, J., Xu, K., Wood, J., Nguyen, P. H., Stephens, D., Wong, B., and Zheng, Y.: *An extensible framework for provenance in human terrain visual analytics*. In *Visualization and Computer Graphics, IEEE Transactions on*, volume 19(12):2139–2148, 2013.

[10] Barthet, M., Plumbley, M., Kachkaev, A., Dykes, J., Wolff, D., Weyde, T. Big Chord Data Extraction and Mining. Proceedings of Conference on Interdisciplinary Musicology, Berlin, 2014.