Exploring Real-time Visualisations to Support Chord Learning with a Large Music Collection

Johan Pauwels
Centre for Digital Music
Queen Mary University of London, UK
j.pauwels@qmul.ac.uk

Anna Xambó
Centre for Digital Music
Queen Mary University of London, UK
a.xambo@qmul.ac.uk

Gerard Roma
CeReNeM
University of Huddersfield, UK
g.roma@hud.ac.uk

Mathieu Barthet
Centre for Digital Music
Queen Mary University of London, UK
m.barthet@qmul.ac.uk

György Fazekas
Centre for Digital Music
Queen Mary University of London, UK
g.fazekas@qmul.ac.uk

ABSTRACT

A common problem in music education is finding varied and engaging material that is suitable for practising a specific musical concept or technique. At the same time, a number of large music collections are available under a Creative Commons (CC) licence (e.g. Jamendo, ccMixter), but their potential is largely untapped because of the relative obscurity of their content. In this paper, we present *Jam with Jamendo*, a web application that allows novice and expert learners of musical instruments to query songs by chord content from a large music collection, and practise the chords present in the retrieved songs by playing along. Its goal is twofold: the learners get a larger variety of practice material, while the artists receive increased exposure. We experimented with two visualisation modes. The first is a linear visualisation based on a moving time axis, the second is a circular visualisation inspired by the chromatic circle. We conducted a small-scale thinking-aloud user study with seven participants based on a hands-on practice with the web app. Through this pilot study, we obtained a qualitative understanding of the potentials and challenges of each visualisation, which will be used to inform the next design iteration of the web app.

1. INTRODUCTION

Repeated practice is an intrinsic part of music education, but can also be monotonous and disengaging. Typically, music students are introduced to new techniques in class, and need to further practise at home by repeating a small number of pieces containing the new material. The reason for this limited number of pieces can be purely practical: it is hard for teachers to find good practice material that fits the class and is musically interesting at the same time. Meanwhile, an increasing number of large music collections are released under a Creative Commons (CC) licence (e.g. Jamendo, ccMixter). These large databases have a great potential because they offer music that is free to use, but they generally reach only a particular audience.

We have developed a system named *Jam with Jamendo*, which attempts to bring the communities of music students and CC artists together for their mutual benefit. We focus on a scenario where music learners are querying for songs based on their chord content. They are invited to select the chords they want to practise and submit their selection to the system, which returns a list of songs for them to play along with. This way, music learners can increase their amount of practice material, while artists gain additional exposure. In this scenario, users can improve their skills in a variety of ways. For instance, learners can try to play on their instrument the chords of the music piece as displayed, either as a sight-reading exercise or working slowly and repeatedly through them, and practise chords transitions along the way. Alternatively, their improvisation skills can be improved by coming up with melodies that fit the chord sequence or substituting chords for alternative ones.

We have automatically extracted chord sequences for 100 000 songs of the Jamendo music database to create a database that can be queried by chords. In order to facilitate playing along with the suggested songs, these chord transcriptions need to be presented to the user in sync with the music.

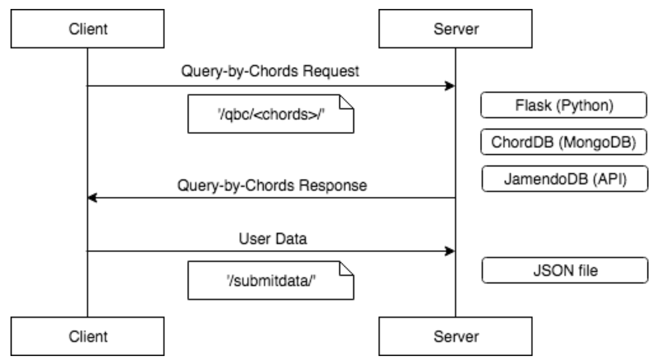

Figure 1: Architecture of the system.

Web Audio Conference WAC-2018, September 19–21, 2018, Berlin, Germany.

The main topic of this paper is the exploration of two different visualisation modes for the chord sequences: (1) A linear visualisation, based on a moving time axis. (2) A circular visualisation, inspired by the chromatic circle. We conducted a small-scale thinking-aloud user study with seven participants based on a hands-on practice with the web app. It allowed us to obtain a qualitative understanding of the potentials and challenges of each visualisation, which will inform the next design iteration of the web app. Our preliminary findings point to the potential of providing two visualisation modes and how to refine the design to satisfy the music learners' needs. This paper is a follow-up of our previous work [21], where we focused on the retrieval scenario and the evaluation of the quality of the songs provided.

2. BACKGROUND

2.1 Information Visualisation

Information visualisation (IV) is an interdisciplinary field that started at the beginning of the 1990s as a collection of techniques that allowed a user to represent and interact with complex data sets [6, 7, 18]. Its main goal has been to understand complex structures of information, showing inner structures and dynamics of change. IV techniques have been applied to areas such as genetics or economics [7]. There is a close link between information visualisation and information retrieval (IR), given that the latter needs to offer alternative ways of searching and retrieving large databases using visual-spatial strategies [2]. Making music visible has been a topic of research where varied strategies of music representation have been explored, such as evoked images, time or frequency content of the audio signal or score representations [9]. In particular, visual metaphors have been used for teaching music theory concepts [18].

2.2 Visualisation in MIR

Music in digital format has facilitated an increase in the size of personal audio collections given the ease to share and download the items as well as the growing storage capacity of both portable and desktop devices. The visualisation of large data collections of audio has been a common topic of research within the area of music information retrieval (MIR), as discussed in [15]. Bidimensional or tridimensional representations of the entire collection have been used according to the concept of similarity between the musical elements, with mappings from a multidimensional to a two- or three-dimensional space. Principal component analysis (PCA) [20] or Self-Organizing Maps (SOM) [19] are well-known techniques for that purpose. The visualisation of single sounds as 2-dimensional signals has been broadly explored (e.g. spectrograms, waveforms, similarity matrices), which is helpful for getting a better sense of the audio content and the properties of a given audio sample [8, 12].

2.3 Chord Visualisation

The graphical representation of chords for a better understanding of the relationships between their comprising notes has been used as a method for teaching music since the days of Pythagoras (6BC). An early influential concept is the theory of the "circle of fifths", where a circle is divided into 12 equal parts, a model further developed as the 'Spiral' by Descartes, as reported in [1]. A more modern graphical representation is Lerdahl's "tonal pitch space" [13].

A number of online systems for music learning based on chords exists, however they tend to be song-centric, meaning that first users decide what song they want to play and then get the chords for it. Thus, music discovery is missing, in contrast to our chord-centric approach. An exception is Hotttabs [3], but the discovery here is based on song popularity and overall difficulty, without possibility to query for the actual chord content. Furthermore, it is guitar-specific, while our system is instrument-agnostic.

Some of the song-centric chord learning systems rely on automatic chord recognition. For example, both Riffstation[1] and Chordify [10] are web and mobile applications that transcribe chord labels for songs retrieved from large online collections such as YouTube or user's personal collections. The chord labels are then visualised in sync with the music. Chordify uses an interface that is reminiscent of a paper chord sheet, with chord symbols delineated by beats and barlines, and a moving cursor that indicates the current position in the song. In contrast, Riffstation has a fixed cursor, and chords are represented as coloured blocks that scroll by the cursor. The information provided includes the current chord as a highlighted block as well as past and future chords which are shown in a dimmer colour. The width of these blocks is proportional to the duration of the chords.

Other visualisations of chord sequences can be found in Yousician and Songle. Yousician[2] is a music learning platform, focussing on guitar, piano, bass and ukulele. It gives real-time feedback on the accuracy and timing of your playing, through a scrolling representation of a virtual fretboard or keyboard. Songle [11] provides a real time visualisation of the structure and content of a song including beat, melody and chords. The chord visualisation is only a small part of a large display, and as a result, the overall interface of this system is quite complex. In this work, we compare two visualisation modes for the playback task that are instrument-agnostic, by allowing users to switch between them and pick the one that best fulfils their needs.

3. SYSTEM

3.1 Concept

As shown in Figure 1, *Jam with Jamendo* is a client-server web application based on a query-by-chords (QbC) approach. The system is designed for supporting music learning with a musical instrument (e.g. guitar, piano, bass guitar, harp). The workflow is based on letting users pick a set of chords that they know and want to practise. Based on this QbC request, a curated list of songs from the Jamendo music collection is provided. Real-time information about the chords within the song is shown, so that music learners can follow along with their instrument. Compared to more traditional approaches to music learning with printed musical scores, this approach allows the user to discover new bands that are licensed under Creative Commons. It also provides multiple ways of representing chord information, which can be especially helpful for the music learner. This work follows up previous research where we explored the retrieval scenario and the quality of the songs listed [21], whereas here we focus on the playback scenario that includes two chord visualisation modes, a linear mode and a circular mode.

[1] https://www.riffstation.com
[2] https://yousician.com

3.2 Back-end

The back-end of the system consists of a MongoDB, a type of document-oriented database based on JSON. It is populated with automatic chord transcriptions of nearly 100 000 songs of the Jamendo Licensing[3] catalogue. Jamendo Licensing is an opt-in service offered by the Jamendo music platform that aims to license out the content of users who signed up to it for commercial purposes. Examples of common use-cases are as part of in-store radios or as background music for video and games. The audio and metadata of these music pieces is accessible through an API.[4]

Because this catalogue is offered to commercial partners, it is lightly curated for recording quality. This means that the team of Jamendo selects tracks of sufficient technical standards to actively promote, without judging their artistic merits. Bad recordings, copyrighted and joke submissions are therefore removed. Of the original 200 000 tracks that are present in the Jamendo Licensing catalogue, about half are withheld by this process. At the time of creation, we managed to get the audio for 99 960 of them through the API. These files were processed off-line by the chord estimation algorithm presented in [16]. A unique property of this algorithm is that in addition to the estimated chord sequence, it also returns a confidence per file, which indicates how sure the algorithm is of its output. This measure will be used to improve the user experience by suggesting those songs of which the chord transcription quality is deemed high first.

We settled on a chord vocabulary of 60 chords: five chord types for each of the 12 possible roots. The chord types are major, minor, dominant 7th, major 7th and minor 7th, which were selected because they came up as the most popular chord types in a study of popular music and together they provide a nearly complete coverage for most songs [5]. The user's selected chords are retrieved by the front-end and are used to suggest a list of songs that contain them. "Songs containing the chords" can be interpreted in many ways, but given the learner's scenario, we decided to return songs in which only the given chords are used. The reasoning is that we want to avoid that the users are presented with unrequested chords, which they possibly do not know how to play. A consequence is that the more chords are selected, the larger the subset of songs that can be suggested. At the same time, we deem it unsatisfactory if a large number of chords are selected and songs get returned that contain only a small number of distinct chords. Therefore preference is given to songs that contain as many as possible of the specified chords (but strictly excluding unspecified chords). In order to deal with the inevitable errors in chord output due to the automatic transcription system, the suggested songs are sorted in decreasing order of confidence in the estimated chord sequence.

3.3 Front-end

The front-end of the app has been built based on a model-view-controller approach. The technologies used include Python's Flask[5] (model), Bootstrap[6] (view) and jQuery[7] and JavaScript (controller). As Figure 1 illustrates, the model (Flask's server) manages the QbC requests and communicates with the databases ChordDB (chord sequences and confidence) and JamendoDB (audio and metadata) to retrieve a list of songs that contain the requested chords, as explained in section §3.2. The view implements the query and the playback interfaces. In the QbC interface, users can select a set of chords based on their knowledge from a provided set of 60 chords. The playback interface shows a list of selected songs and two chord visualisation modes, a linear view and a circular view. The controller manages the user's interaction with the GUI and secures the asynchronous communication with the Jamendo API to stream the songs. A companion webpage with extra information and a demo video can be found online.[8]

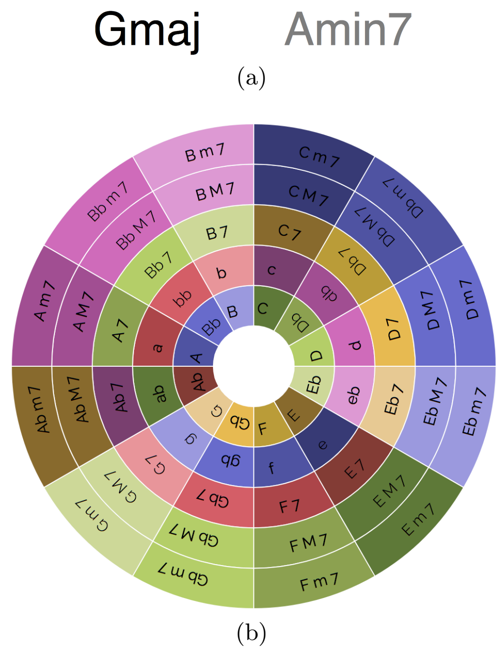

Figure 2: (a) Screenshot of the linear visualisation. (2) Screenshot of the circular visualisation.

4. VISUALISATIONS

Traditionally, visualisations of single sounds are based on the perceptual domain, such as spectrograms [12, 15]. Here we propose to explore an intuitive, cognitive-perceptual model that allows the user to hear by seeing the chords. We provide two visualisations: (1) Textual information about the current and future chord only. (2) Textual information about the current and future chord within the context of a chromatic circle of chords.

4.1 Linear Visualisation

We implemented a linear visualisation inspired by the concept of a moving time axis. This approach is similar to the visualisation used in the Riffstation web application, which contains a train of moving chords. Our approach is stripped-down, more space-efficient version that focusses only on showing the present chord in black text and the next chord in grey text, as shown in Figure 2a.

[3]https://licensing.jamendo.com
[4]https://developer.jamendo.com
[5]http://flask.pocoo.org
[6]https://getbootstrap.com
[7]https://jquery.com
[8]https://www.audiocommons.org/jam

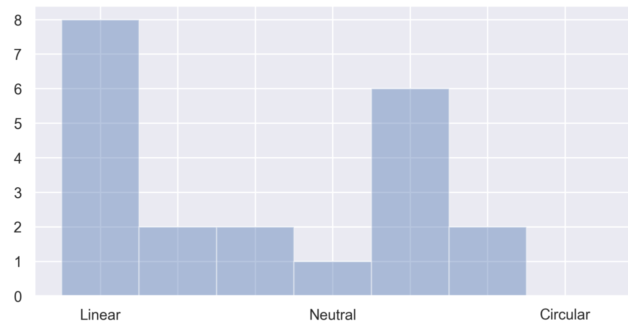

Figure 3: Bar plot of the ratings of the two visualisation modes ($N = 21$).

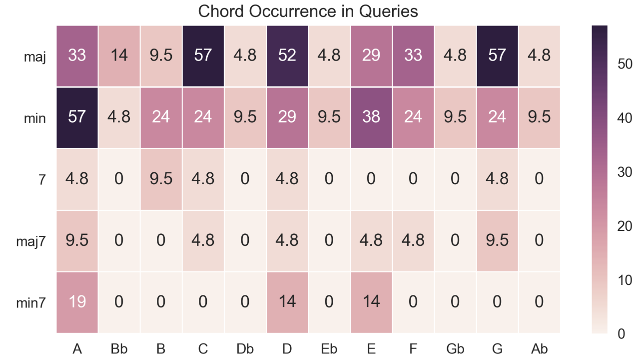

Figure 4: Proportion of the queries ($N = 21$) that contain each chord.

4.2 Circular Visualisation

As illustrated in Figure 2b, we developed a circular visualisation inspired by the circular representation of chords. In particular, we implemented a chromatic circle, closely related to the circle of fifths, as discussed in section §2.3. We used the web-based JavaScript library D3 [4], which is designed to create interactive data visualisations. In particular we used the template for D3 sunburst visualisation (partition layout).[9] The colours were arbitrarily distributed using an ordinal colour scale under the *schemeCategory20b* option.[10]

5. USER STUDY DESIGN

An experiment was designed so that we could observe either in person or remotely the participants interacting with the system while practising with their musical instrument. With each of the participants, we had a conversation about their opinions on the visualisation modes inspired by thinking-aloud techniques [17]. We captured their thoughts in real-time by taking notes as a form of verbal protocol.

The participants were asked to bring their musical instrument or use a provided electric guitar and (1) choose a set of chords that they know from the web application; (2) explore the suggested three songs containing the chords by using the two visualisation modes and playing along; (3) rate their preferred visualisation mode from a 7-point Likert scale ranging from "strongly prefer linear" (StP1), "prefer linear" (P1), "slightly prefer linear" (SlP1), "neutral" (N), "slightly prefer circular" (SlP2), "prefer circular" (P2) and "strongly prefer circular" (StP2). They were asked to repeat this task three times in independent trials. The task was expected to be completed in 45—60 minutes. A final post-questionnaire gathered demographic information.

6. FINDINGS

We obtained responses from seven participants (four female, three male) with an age range of 26–43 years old (M=33.86, SD=7.06). In the context of usability, seven users has been found to be a suitable number for a small-scale user study [14]. We present our findings as a preliminary study that is helpful to identify potential challenges of the prototype and design the next iteration. From here on we will use the nomenclature of P1–P7. We thus gathered 21 independent trials. Three of the participants did the study remotely while the other four did the study co-located. The participants' musical skill level was intermediate (3 participants) and advanced (4 participants). Their experience with music technology varied from less than 2 years (1 participant), 4–6 years of experience (1 participant), to more than 10 years of experience (5 participants). The instruments used for the study included electric guitar (3 participants), piano (2 participants), Roli Seabord (1 participant) and a synth (1 participant).

Figure 3 shows the aggregation of the 7-point Likert scale regarding the preferred visualisation mode. The linear visualisation was the most popular: on 12 occasions the linear visualisation was preferred, 8 of which were a strong preference. In contrast, there were 8 occurrences in which the circular visualisation was preferred, 6 of those being a slight preference. The number of chords selected by the participants ranged from 1 to 24, with an average of 6.81 ± 5.74. The queries are visualised in Figure 4 as the proportion of total queries that contain each chord. We can see that there was a strong preference for major and minor chords, with *Cmaj*, *Gmaj* and *Amin* as the most popular ones. In general, the most popular roots were natural notes. This preference largely reflects the chords that are taught first in music curricula, and therefore should not come as a surprise.

Tables 1 and 2 summarise the findings from the thinking-aloud conversations with the participants. From Table 1 we observe that even though there was a stronger preference for the linear visualisation, there were arguments in favour of both. The linear visualisation was associated with simplicity and the circular visualisation tended to be associated with more complexity. The use of colours seemed to be more appealing than using black-and-white text, however they should have a clear meaning. For example, the relationships between chords and chord types, or the specified vs. unspecified chords, could be chromatically reflected. Table 2 gives an overview of the main aspects that could be improved from the web application: chord recognition, QbC requests, context of the song and practice hints. Overall, the participants found the application to look promising, but a number of them suggested it should blend better into the music learning practice and the real-time activity of playing a song.

[9] https://bl.ocks.org/denjn5/f059c1f78f9c39d922b1c208815d18af
[10] https://bl.ocks.org/pstuffa/3393ff2711a53975040077b7453781a9

Table 1: Benefits and Challenges Between Linear and Circular Visualisations

Type	Benefits	Challenges
Circular	• The fixed position of the chords (2 participants): *"no need to read it, less effort"* (P2) and *"you can see all the chords"* (P6). • The relationships between chords (2 participants): it promotes seeing the *"harmony, patterns, (...) and finding relationships"* (P4). • The colours (3 participants): *"it is more pleasing aesthetically"* (P7). • Intuitive (1 participant): *"the colour can be more intuitive for playing the chord by colour and position"* (P6) and *"it provides more information and can be remembered in a more intuitive way"* (P6). • Usage (2 participants): *"if you can follow the song easily because you know the chords, then you can spend sometime looking at the circle and finding relationships"* (P4) and *"it's fun now that I am more confident to play with the chords."* (P4).	• Missing a global structure (3 participants): *"I know what the next chord is, but I don't know when it is going to change"* (P6). • Unclear relationships (1 participant): *"Make the circle turn such that the tonic is in a fixed position based on key recognition"* (P4). • Readability (4 participants): *"it is hard to read, there is too much information"* (P1) and *"unsettling colours"* (P3). • Complexity (4 participants): *"it is difficult to see what is happening"* (P3) and *"my knowledge of music theory does not give me a clear understanding of the colour wheel"* (P7). • Usage (1 participant): *"if the chords change fast they are hard to follow and need to pay more attention than in the timeline"* (P4) and *"I would not use it for a concert"* (P4).
Linear	• Simplicity (2 participants): *"I don't feel confident with the chords, so it is easier to see the current and next chord (no distraction)"* (P4). • Usage (2 participants): *"I find this very useful as an improviser"* (P3) and *"this works well if you need to pay a lot of attention, for example, if the chords change fast"* (P4). • Clarity (2 participants): *"it is faster and clearer with the letters"* (P7).	• Missing a global structure (3 participants): *"I know what the next chord is, but I don't know when it is going to change"* (P6) and *"for songs that don't change chords every 4 beats it is harder, there are a lot of transitions that kill you"* (P4). • Lack of relationships (1 participant): *"seeing the connections would be nice (...) for solos see substitutions (that's a problem that I have when learning), for example when doing arrangements"* (P3). • Lack of colours (1 participant): *"there could be more colours"* (P1).

Table 2: Design Issues and Design Suggestions for Improvement

Themes	Design Issues	Design Suggestions
Transcription	• Chords are delayed and sometimes the chord transcription is not accurate (4 participants): *"chord delays confuses me especially when they change fast"* (P2) and *"are the chords correct? they don't change aligned to the music (there's a little bit of a lag)"* (P4).	• Show the information to allow for anticipation (1 participant) *"Need to be able to anticipate"* (P1) and combine automatic chord recognition with other MIR techniques, e.g. *"combine symbolic recognition to improve the chord recognition with the confidence estimation"* (P1).
QbC	• The QbC interface could align better with music education (3 participants): *"it's better now that I reduced the number of chords and I could focus more on playing"* (P4); *"querying by tonality and adding diminished chords would be useful"* (P5) and *"the chord structure is not the way I use for classical piano study, but scales with their accidentals. It's then difficult to understand and reproduce the songs following the chords chart."* (P7).	• Reconsider the organization the QbC grid (1 participant): *"the chords could be organized by notes e.g. C major, C minor, C 7, and so on, for both expert musicians who are familiar with chords, music solfège, harmony, but also for the novice musicians"* (P5).
Song	• If the song is unknown by the practitioner, it is difficult to identify when the next chord comes and therefore to practise (4 participants): *"Need to know about the musical structure e.g. chorus, verse"* (P1); *"it would be nice to have a sense of duration beside the chord, how long it is going to last (to allow for anticipation strategies when performing)"* (P3) and *"when rehearsing with the system you are never told what octave are you in"* (P3).	• Show more contextual information about change and temporal durations (4 participants): *"include beat timing"* (P1) and *"it would be nice to see when the change is coming"* (P3).
Hints	• It is unclear how many chords are suitable to include in the QbC request (the 7 participants asked how many chords they should query): *"now I've selected a more complex song because it could be more interesting"* (P6) and *"it's not so useful when there is only one chord"* (P6).	• Give some (automatic) hints on the types of chords to choose or play (3 participants): *"exploring chords that you know with one or two new chords could be interesting"* (P1) and *"with this kind of music (...) it makes more sense to improvise"* (P3).

7. LESSONS LEARNED

From this experience, we have identified four main areas that we would like to improve in the next design iteration of *Jam with Jamendo*:

- **Improvement of the user experience**: address the delay between current and next chord and provide more contextual information of the song (e.g. representation of metrical structure) to facilitate sight-reading.

- **Keep two visualisations**: it seems useful to provide two modes of visualisation so that users can pick according to their musical knowledge, personal taste and needs. The use of colours seems to be useful, yet a colour palette needs be explored more thoroughly for both visualisations.

- **Improvement of the QbC interface**: it might be interesting to suggest chords that go well with the user's selection based on occurrence in the dataset and also suggest alternative chords informed by music theory.

- **Expanding the features of the web application based on real use-cases**: from this user study a number of usages emerged, such as music improvisation or teaching, which can inform the new forthcoming features of the application in line with co-design practices.

8. CONCLUSION

In this paper, we presented *Jam with Jamendo*, a web application designed for learners of musical instruments that supports query-by-chord requests from the music database Jamendo. In particular, we conducted a thinking-aloud user study with seven participants that compared two visualisation modes: a linear mode and a circular mode. We identified the benefits and challenges of each visualisation and a set of design issues and potential solutions for the next iteration of the web application. In summary, using two visualisation modes is promising, but they need to better support the task of real-time performance within an educational context.

As future work, we plan to improve the existing visualisations, add more contextual information of the metric structure, and also improve the query-by-chord task. So far, we have relied on the number and type of chords to implicitly infer the difficulty of music pieces, but it would be useful to present the user with the option to exclude or order pieces according to difficulty. After all, the same chords can be used in a wide range of complex rhythmical patterns. Analysing the rate of chord changes would be a first step towards a solution.

9. ACKNOWLEDGMENTS

We are grateful to the participants of the study. AudioCommons is funded by the European Commission through the Horizon 2020 programme, research and innovation grant 688382. This work has been partly funded by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/L019981/1.

10. REFERENCES

[1] A. Abraham. The Zone of Musical Creativity: Harmonic Series Structure – From Pictorial Representation to a Method of Teaching. In *Proceedings of the 7th International Conference on Knowledge, Information and Creativity Support Systems*, pages 207–210, 2012.

[2] R. A. Baeza-Yates and B. A. Ribeiro-Neto. *Modern Information Retrieval*. ACM Press/Addison-Wesley, 1999.

[3] M. Barthet, A. Anglade, G. Fazekas, S. Kolozali, and R. Macrae. Music Recommendation for Music Learning: Hotttabs, a Multimedia Guitar Tutor. In *Workshop on Music Recommendation and Discovery*, pages 7–13, 2011.

[4] M. Bostock, V. Ogievetsky, and J. Heer. D^3 Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.

[5] J. A. Burgoyne, J. Wild, and I. Fujinaga. An Expert Ground-Truth Set for Audio Chord Recognition and Music Analysis. In *Proceedings of the 12th ISMIR Conference*, pages 633–638, 2011.

[6] S. K. Card. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, 1999.

[7] C. Chen. *Information Visualization: Beyond the Horizon*. London: Springer, 2nd edition, 2006.

[8] N. Cook. Methods for Analysing Recordings. *The Cambridge Companion to Recorded Music*, pages 221–245, 2009.

[9] M. Cooper, J. Foote, E. Pampalk, and G. Tzanetakis. Visualization in Audio-Based Music Information Retrieval. *Computer Music Journal*, 30(2):42–62, 2006.

[10] B. de Haas, J. P. Magalhaes, D. T. Heggeler, G. Bekenkamp, and T. Ruizendaal. Chordify: Chord Transcription for the Masses. In *Proceedings of the 13th ISMIR Conference Late Breaking and Demo Session*, 2014.

[11] M. Goto, K. Yoshii, H. Fujihara, M. Mauch, and T. Nakano. Songle: A Web Service for Active Music Listening Improved by User Contributions. In *Proceedings of the 12th ISMIR Conference*, pages 311–316, 2011.

[12] E. Isaacson. What You See Is What You Get: On Visualizing Music. In *Proceedings of the 6th ISMIR Conference*, pages 389–395, 2005.

[13] F. Lerdahl. *Tonal Pitch Space*. Oxford University Press, 2001.

[14] J. Nielsen and T. K. Landauer. A Mathematical Model of the Finding of Usability Problems. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, pages 206–213, 1993.

[15] N. Orio. Music Retrieval: A Tutorial and Review. *Foundations and Trends® in Information Retrieval*, 1(1):1–90, 2006.

[16] J. Pauwels, K. O'Hanlon, G. Fazekas, and M. B. Sandler. Confidence Measures and Their Applications in Music Labelling Systems Based on Hidden Markov Models. In *Proceedings of the 18th ISMIR Conference*, pages 279–285, 2017.

[17] Y. Rogers, H. Sharp, and J. Preece. *Interaction Design: Beyond Human-Computer Interaction*. John Wiley & Sons, 2011.

[18] R. Spence. *Information Visualization*. ACM Press, 2001.

[19] R. Typke, F. Wiering, and R. C. Veltkamp. A Survey of Music Information Retrieval Systems. In *Proceedings of the 6th ISMIR Conference*, pages 153–160, 2005.

[20] G. Tzanetakis and P. Cook. 3D Graphics Tools for Sound Collections. In *Proceedings of the Conference on Digital Audio Effects*, pages 7–9, 2000.

[21] A. Xambó, J. Pauwels, G. Roma, M. Barthet, and G. Fazekas. Jam with Jamendo: Querying a Large Music Collection by Chords from a Learner's Perspective. In *Proceedings of the 13th International Audio Mostly Conference*, 2018.