

Commentary on Shaffer et al.: A cluster analysis of harmony in the McGill Billboard dataset

DANIEL MÜLLENSIEFEN[1]

Department of Psychology, Goldsmiths, University of London

This short commentary on the target paper by “A cluster analysis of harmony in the McGill Billboard” by Shaffer et al. starts with observing that not all harmonic progressions that are theoretically possible are equally common. Instead, some progressions are more popular than others in popular. In fact, certain harmonic progressions are closely associated with specific styles and sub-genres and it is the aim of the target paper to provide a meaningful classification system for harmonic progression. The commentary identifies several strengths of the target paper, including a nice balance between rigorous empirical work and providing a context and interpretations that are musicologically well-informed. In its critique the commentary points to the limitations of only using harmonic bigrams (i.e. the transitions between two chords) as the empirical data and the missing link to related literature on harmonic modelling in the music information retrieval community.

Submitted 2020 May 7; accepted 2020 May 8.

Published 2020 July 6; <https://doi.org/10.18061/emr.v14i3-4.7657>

KEYWORDS: *Chord progressions, harmonic modelling, cluster analysis, popular music corpus*

ONE of the hallmarks of harmony in popular music is that it is very formulaic, i.e. many pop songs share the same sequences of chords. Sometimes harmonic sequences are re-used in varied form, but sometimes repeated use occurs quite literally. The repetition of harmony can certainly contribute to the popularity of a genre because it enables listeners to recognise and process a chord sequence easily and quickly. This ease of processing can contribute to appreciation and increased liking via psychological mechanisms such as the mere exposure effect (Zajonc, 1968). Similarly, the use of common harmonic progressions in pop songs helps musicians to play, sing along and memorise songs. The degree to which chord progressions are re-used across songs varies between different sub-genres of pop music. For blues and rock’n’roll just a few harmonic sequences are necessary to cover the harmonic content of the vast majority of songs. In other genres the variety of harmonic progressions is broader and less central to the definition and recognition of a genre. Yet, even in these subgenres of pop music the use of common harmonic sequences is frequent and harmonic formulae can be found across many songs. Artistic evidence for the ubiquity of certain chord progression provides a famous compilation of fragments from hit songs called “4 Chords” that the comedy act Axis of Awesome has concatenated in a continuous stream [2].

From the perspective of combinatorics there is no need to repeat chord sequences at all. Assuming that we have 12 basic chords at our disposal in a given key, there are $12^8 = 429,981,696$ different possible chord sequences for an eight-bar song fragment if chords were to change for each bar. But only very few of these chord sequences are suitable for the use in pop music. Popular music is popular because it makes use of musical (here more specifically: harmonic) material that has been used before and that listeners like. This explains why only a very few of the 429,981,696 possible eight-bar sequences occur at all in commercial pop music, and why some of these are much more frequent in pop music than others.

THE TARGET PAPER

The very unequal distribution of harmonic sequences in popular music provides an opportunity to discover relationships between songs by identifying shared chord progressions. The more chord progressions are shared in a set of songs, the more they can be considered to be harmonically similar. Thus, grouping



harmonically similar songs into clusters can provide a potentially very interesting and useful classification and description system that can be used to trace stylistic relationships and understand “the grammars and practices” of song writing in popular music. This is the main goal of the harmonic cluster analysis that Kris Shaffer and colleagues perform in their article “A cluster analysis of harmony in the McGill Billboard dataset”. To obtain an empirical measurement of harmonic content Shaffer and colleagues determine transition probabilities between chords, i.e. they tally the number of instances where chord x is followed by chord y . They measure harmonic relationships in each of the 730 songs of the McGill Billboard corpus by the transition frequencies of all 144 theoretically possible chord transitions. Then, Shaffer et al. use cluster analysis to put those songs together into the same cluster that share a similar profile of transition probabilities. However, the k-means clustering algorithm they use does not offer an in-built solution for the number of clusters to choose. This is where Shaffer et al. draw on different sources from pop music analysis and previous work on genre classification. Based on the notion of the “six tonal grammars” of rock music postulated by rock analyst Walter Everett (2004) they analyze clustering solutions from 1 to 8 clusters, where different clusters comprise songs with different harmonic “grammars and practices”.

STRENGTHS

One of the strengths of the article by Shaffer and colleagues is that they resist the temptation of wanting to discover the “true” number of distinguishable harmonic grammars in pop music. Instead, they provide a compact and intelligently labelled map of emergent clusters of harmonic relationships (their Figure 2) that ranges from the single cluster containing all songs of the corpus to the much more refined classification of the 8-cluster solution. Most helpfully, they also provide a description and interpretation of each cluster within all eight clustering solution, including many song examples and references to other markers or musical style and genre. All cluster interpretations are supported by graphs depicting the most frequent chord transitions found in this cluster. These very clear and knowledgeable interpretations help the reader to make sense of the large amount of information that these cluster solutions comprise. The clusters that Shaffer and colleagues present represent sets of harmonic progressions that frequently occur together in the songs of a cluster—or common “harmonic practices” to use their term. Because these harmonic practices are derived empirically, they are a very useful data-based complement to manual analyses of harmonic formulae in pop that are often based on small samples. In fact, I believe that the empirical data that Shaffer and colleagues present can be a very useful starting point for subsequent qualitative investigations of the use harmonic formulae in individual songs. Only with the empirical frame of which chord transitions are common in certain pop styles can we begin to understand what the deviations and extensions of these common patterns in individual songs mean.

CRITIQUE

However, there are a few areas where the article by Shaffer and colleagues is limited and leaves some open questions that may, however, be answered in future research. The main disciplinary background of their article is the musicology and analysis of popular music. This firm background in music analysis is one of the strengths of their article as it informs their hypotheses (e.g. on the possible number of clusters) and their interpretation of the empirically derived clusters. But at the same time their summary of the literature in the introduction lacks an acknowledgement of relevant computational studies on chord progressions that have been carried out in the field of music information retrieval (MIR). Conceptually related is for example the study by Mauch et al. (2007) on common chord progressions in jazz and Beatles corpora. Another related paper was published by Scholz et al. (2009) and shows how chord sequences can be modeled by n-gram techniques. In addition, there are a few studies that present related computational-musicological work on chord sequences extracted from pop music corpora. Among these studies is the paper on big chord data mining by Barthet et al. (2014) and the two different data-driven studies of the evolution of popular music by Serra et al (2010) and Mauch et al. (2015) which also make use of harmonic descriptors. A closer look at these related approaches from MIR might have helped to refine and enrich the interpretations of the harmonic clusters in the article by Shaffer and colleagues.

A second point of criticism concerns the use of their criteria “nuance” and “generalizability” to evaluate cluster solutions. Their general point is clear: A convincing solution should achieve a balance between a nuanced solution, distinguishing between songs with subtle but meaningful differences in their harmonic patterns, and a general solution, classifying songs with similar features into the same cluster. But

then the comparative evaluation of the eight different clustering solution only seems to make superficial use of these two criteria.

Finally, the empirical approach presented by Shaffer et al. is limited by relying only on harmonic bigrams, i.e. only considering the transition probability between two chords. It is tempting to interpret the sets of directed arrows in their cluster graphs as actual sequences of four, five or six chords. But in fact, from these cluster graphs we cannot tell whether one pair of chords is frequently followed or preceded by another pair of chords or whether these two pairs occur in succession at all in the songs of a cluster. Without a doubt, the transitional probabilities from one chord to the next represent meaningful harmonic information. But harmonic formulae or idiomatic chord progressions are usually considered to be sequences of three or more chords. To capture longer progressions, the harmonic bigram approach can be extended to include sequences of varying lengths, giving rise to n-gram approaches as explored for example by Scholz et al. (2009). Hence, perhaps the time is ripe now for exploring longer chord progressions and harmonic formulae in the McGill Billboard corpus.

In conclusion, the target paper by Shaffer and colleagues provides a very valuable contribution to the empirical study of harmony in popular commercial music. It combines perspectives from popular music analysis and computational musicology, and researchers with backgrounds in either of these fields will discover interesting ideas and results in this paper. From the wider perspectives of these two areas some limitations of the paper are noticeable. However, I am certain that future empirical studies on harmony in popular music will build on the results and the inspirations of this paper.

ACKNOWLEDGEMENTS

Parts of this commentary build on thoughts on the target paper provided by Matthias Mauch. This article was copyedited by Daniel Shanahan and layout edited by Diana Kayser.

NOTES

[1] Correspondence can be addressed to: Prof. Dr. Daniel Müllensiefen, Department of Psychology, Goldsmiths, University of London, New Cross Road, New Cross, SE14 6NW, d.mullensiefen@gold.ac.uk.

[2] Axis of Awesome. 4 Chords. Video: <https://youtu.be/oOIdewpCfZQ>

REFERENCES

- Barthet, M., Plumbley, M. D., Kachkaev, A., Dykes, J., Wolff, D. and Weyde, T. (2014). Big Chord Data Extraction and Mining. Paper presented at the *9th Conference on Interdisciplinary Musicology – CIM14*, Staatliches Institut für Musikforschung, Berlin, Germany. Retrieved from <http://openaccess.city.ac.uk/5803/>.
- Everett, W. (2004). Making Sense of Rock's Tonal Systems. *Music Theory Online*, 10 (4). Retrieved from http://www.mtosmt.org/issues/mto.04.10.4/mto.04.10.4.w_everett.html.
- Mauch, M., Dixon, S., Harte, C., Fields, B. & Casey, M. (2007). Discovering Chord Idioms Through Beatles and Real Book Songs. In: *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007)*. Vienna, Austria.
- Mauch, M., MacCallum, R.M., Levy, M., Leroi, A.M. (2015). The evolution of popular music: USA 1960–2010. *Royal Society Open Science* 2, 150081. <https://doi.org/10.1098/rsos.150081>
- Shaffer, K., Vasiente, E., Jacquez, B., Davis, A., Escalante, D., Hicks, C., McCann, J., Noufi, C., Salminen, P. (2019) A cluster analysis of harmony in the McGill Billboard dataset. *Empirical Musicology Review*. Vol.14, no.3-4. 146-162. <https://doi.org/10.18061/emr.v14i3-4.5576>

Scholz, R., Vincent, E., & Bimbot, F. (2009). Robust modeling of musical chord sequences using probabilistic N-grams. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2009), Taipei, (pp. 53-56). <https://doi.org/10.1109/ICASSP.2009.4959518>

Serrà, J., Corral, Á., Bogaña, M., Haro, M., & Arcos, J.Ll. (2012). Measuring the Evolution of Contemporary Western Popular Music. *Scientific Reports* 2, 521. <https://doi.org/10.1038/srep00521>

Zajonc, R. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology Monograph Supplement*, 9, 1–27. <https://doi.org/10.1037/h0025848>