INTEGRATION OF CROWD-SOURCED CHORD SEQUENCES USING DATA FUSION

Hendrik Vincent Koops, W. Bas de Haas, Anja Volk
Department of Information and Computing Sciences
Utrecht University , the Netherlands
{h.v.koops,w.b.dehaas,a.volk}@uu.nl

ABSTRACT

We propose a novel way to integrate a large amount of crowd-sourced chord sequences with the goal of improving automatic chord extraction. Through Chordify, an online automatic chord extraction service where users can create their own personalized chord sequence, a large amount of different chord label sequences of popular songs are obtained. Using *data fusion* with *copy detection* techniques, we integrate the conflicting chord labels of a song into one, better chord sequence. To evaluate the integrated chord sequences, we measure play time statistics on the Chordify website, and hypothesize that more play time is associated with a chord sequence that matches the audio better. In an initial comparison with a majority vote algorithm and the original automatic chord extraction output, we find that users do prefer data fusion integration over simple majority vote integration.

1. INTEGRATION OF CHORD SEQUENCES

With the recent rapid growth and expansion of web sources containing user-generated content, a large amount of conflicting data can be found in various domains. For example, different encyclopediae can provide conflicting information on the same subject, and reviews and pictures of the same product can differ among websites. An example in the music domain are websites that offer symbolic representations of how to play popular songs, such as tabs or chords. These websites often provide multiple, conflicting versions of the same song, thereby providing different chord label sequences for the same songs. In this research, we address the problem of finding the most appropriate chord labels within a large amount of conflicting chord label sequences.

1.1 Chordify

A recent example of a web source with user-generated content is *Chordify* [2], an online automatic chord extraction (ACE) service that lets users edit the ACE output and create their own personal version of a song. Currently, Chordify is used by 1.5 million musicians every month, and these musicians use the site around 4000 hours every day. Chordify uses the HarmTrace [1] model to extract the chords from the audio. Because ACE is never perfect, Chordify extended their service with an intuitive editing interface allowing users to correct small mistakes. Users can insert, delete, replace, and shift chords. Personalization of the ACE output has resulted in a large amount of conflicting chord sequences, which provides invaluable knowledge with regard to musical variation, music perception and distributed muscial knowledge. However, how these crowd-sourced chord sequence corrections can be used to overcome the transcription errors of ACE remains an open problem that is addressed in this research for the first time.

1.2 Data Fusion

We investigate the problem of integrating a large number of conficting muscial sources using *data fusion* techniques [5]. We refer to different versions of the same song as *sources*, each providing a sequence of chord *values*. See Table 1 for an example. The goal of *data fusion* is to find the *true* values between conflicting sources, meaning: the values with the highest probabilities. We take into account 1) the *accuracy* of sources, 2) the probabilities of the chord labels they provide, and 3) the probability of *copying* between sources.

The accuracy of a source is defined as the average of the probabilities of the chords it provides, and can be seen as the probability that a chord provided by a source is the true chord. In online crowd-sourced data sources (partly) copying information is a common phenomenon. Hence, Dong et al [3, 4] developed a method to detect copied values automatically. Also on Chordify users can copy the chords from another user, this makes it possible that chords are

	T_0	T_1	T_2	T_3	T_4	T_5	T_6
S_0	C:maj	F:maj	G:maj	C:maj	C:maj	G:maj	C:maj
S_1	C:maj	F:maj	G:maj	A:min	A:min	G maj	C:maj
S_2	C:maj	F:maj	C:maj	A:min	C:maj	C:maj	C:maj
S_3	C:maj	C:maj	G:maj	A:min	C:maj	C:maj	C:maj

Table 1. Example of 5 sources ($S_{0...4}$) providing conflicting and identical chords at seven time points ($T_{0...6}$). `maj` and `min` denote a major or minor chord, respectively.

© Hendrik Vincent Koops, W. Bas de Haas, Anja Volk. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Hendrik Vincent Koops, W. Bas de Haas, Anja Volk. "Integration of Crowd-Sourced Chord Sequences using Data Fusion", Extended abstracts for the Late-Breaking Demo Session of the 16th International Society for Music Information Retrieval Conference, 2015.

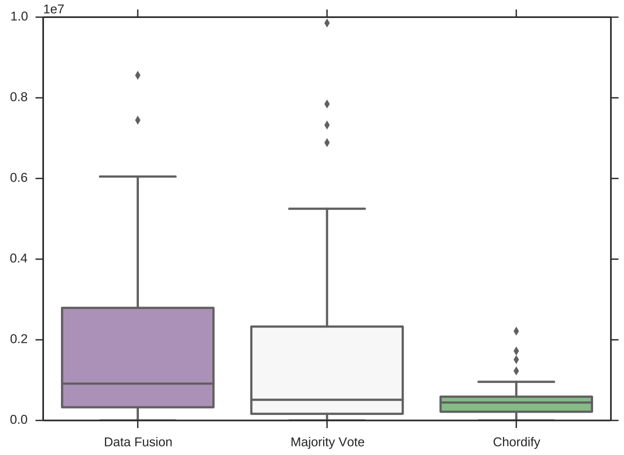

Figure 1. Average playtime of 33 songs in milliseconds.

transitively copied by a large number of other users. Detecting copying between sources aims to correct for this bias. Detecting copying between sources hinges on the intuition that the probability that two independent sources share the same chord label with a low probability is unlikely. In other words, sharing values with low probabilities is evidence for copying, and is used to weight chord labels negatively.

We iteratively computing copy detection, chord probabilities and source accuracies. After convergence of the values, we choose the chords with the highest probabilities, which results in a new chord sequence that integrates the chords from all sources.

For the integration of chord label sequences we employ the ACCUCOPY model that was introduced by Dong et al. in [3, 4].

1.3 Evaluation

The chord sequence improvement is measured in an ecological user setting on the online music e-learning platform Chordify. Chordify synchronizes chords and YouTube videos in order to facilitate learning and playing along with a song. For a large number of songs we measure total play-time on Chordify.net. We assume that if chord sequences better match the YouTube video, it will increase the time they are played on Chordify.

Figure 1 shows the average playtime for the three versions of the songs from a pilot study using 33 songs. We see that the fused chords are indeed are played longer than the original chordify version. In the near future, we plan on repeating this experiment with over 500 songs, which will provide a more detailed view of preferred chord sequences.

2. ACKNOWLEDGEMENTS

H.V. Koops, A. Volk and W.B. de Haas are supported by the Netherlands Organization for Scientific Research, through the NWO-VIDI-grant 276-35-001 to A. Volk.

3. REFERENCES

[1] Bas de Haas, José Pedro Magalhaes, Remco C Veltkamp, and Frans Wiering. Harmtrace: Improving harmonic similarity estimation using functional harmony analysis. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 67–72, 2011.

[2] W Bas De Haas, José Pedro Magalhaes, Dion Ten Heggeler, Gijs Bekenkamp, and Tijmen Ruizendaal. Chordify: Chord transcription for the masses. In *Demonstration presented at the International Society for Music Information Retrieval Conference*, pages 8–12, 2012.

[3] Xin Luna Dong, Laure Berti-Equille, and Divesh Srivastava. Integrating conflicting data: the role of source dependence. *Proceedings of the VLDB Endowment*, 2(1):550–561, 2009.

[4] Xin Luna Dong and Felix Naumann. Data fusion: resolving data conflicts for integration. *Proceedings of the VLDB Endowment*, 2(2):1654–1655, 2009.

[5] Lawrence A Klein. *Sensor and data fusion: a tool for information assessment and decision making*, volume 324. Spie Press Bellinghamˆ eWA WA, 2004.