Query-Based Mashups of Historical Live Music Recordings

Florian Thalmann, Thomas Wilmering, Mark B. Sandler[1]

Navigating and searching large archives of audio or musical material can be tedious due to the temporal nature of the content. Whereas in visual archives one can gain a reasonably reliable overview of search results at a glance simply by scrolling through collated images, with audio material one has to go through search results by listening to individual recordings one by one – a process that may take a considerable amount of time and effort. In this paper we investigate different strategies of automatically creating sound collages based on user-defined queries enabling the user to gain a quick overview of a large number of results. We demonstrate the principles in a prototypical application based on the Grateful Dead collection of the Live Music Archive (LMA)[2].

Background

In recent years, advancements in music information retrieval have led to an improved organization of large digital music collections. In addition to extra-musical metadata, items in musical archives can now be automatically annotated with content-based information ranging from simple audio descriptors to high-level musical analyses. With these annotations one can now satisfyingly query and filter archive content, and quantify or visualize desired characteristics of the results (Elias et al., 2002, Bechhofer et al., 2017). However, for certain tasks it may still take a long time to make sense of these outputs, which can often only be done by listening to the found examples one by one.

To address this problem, we propose the concept of *query-based mashups* which can be generated once the material of an archive is annotated with appropriate audio features. These features can be used to align and reorganize relevant fragments of the material and present them to the users as coherent collages facilitating the exploration of the audio results in an interactive way. Our concept of personalized mashups consists of a three-step process:

- *Selection*: the user designs a search query to find a subset of the database's content. For example, all the recordings of a particular performance, all the versions of one song within a given time span, or all the performances in a particular place. The queries may also contain features automatically extracted from the audio material, in order to select song recordings based on, for example, tempo, timbre, or musical structure.

[1] Centre for Digital Music, Queen Mary University of London
[2] https://archive.org/details/GratefulDead

- *Organization*: the users decide how the resulting material is to be organized or ordered, e.g. by recording date, tempo, or based on content similarity

- *Mashup Parameters*: finally, the users define how the material is mixed together, how much of each result is used (e.g. in number of bars or seconds), what part is used (e.g. the loudest or most varied segment), by how much the segments overlap, and what kinds of transitions are made between them. Part of this process is typically automated and executed by algorithms based on the audio features in the archive.

Demo Application

We implemented a prototypical Web application as part of a platform focusing on live concerts of the band *The Grateful Dead* (Benson, 2016) of which a considerably large number of audience and mixing desk recordings survive in the LMA - over 12,000 from the years 1965 to 1995. The platform allows users to explore the band's concert history in the form of an audiovisual experience by linking the audio recordings with content from several Semantic Web resources such as LMA Linked Data (Bechhofer et al., 2013) or DBpedia[3], as well as data on other Internet resources related to Grateful Dead concerts (Thalmann et al., 2018). These data include setlists and lineup details as well as scans of artefacts such as tickets, posters, photos, or fan mail, stored for instance in the Grateful Dead Archive Online curated at the University of California Santa Cruz (UCSC)[4].

We aggregated data from these sites using dedicated scripts which automatically generate structured RDF data corresponding to our data model based on various combined ontologies (Thalmann et al., 2018). Audio feature extraction results are linked via the Computational Analysis of the Live Music Archive (CALMA) dataset (Bechhofer et al., 2017). In earlier publications, we have shown how Semantic Web technologies are ideal for interlinking various Cultural Heritage resources and searching them in a joint manner (Wilmering et al., 2016, Thalmann et al., 2018).

The main platform allows the exploration of different entities such as shows, venues, locations, songs, tours, or musicians with their aggregated metadata as well as custom playlists, which users can create as they find items on the different pages of the website. In addition to this we experimented with different ways in which the musical material can be explored auditively. One of them automatically mixes the different recordings of one particular concert into an immersive experience by aligning and resampling the recordings in order to match the different tape speeds, and by clustering them into a multidimensional space based on their average distance from each other over time (Wilmering et al., 2016). The user can then interactively explore this space and discover the different qualities and content of the recordings.

[3] https://wiki.dbpedia.org
[4] https://www.gdao.org

In the latest addition, which we demonstrate in our poster presentation, users can choose between different sets of songs, orderings, and parameter settings to create their own personal mashup compositions from the archive material. Mashups are particularly interesting in this context given that one of composer John Oswald's *plunderphonic* albums is based on selected Grateful Dead recordings (Oswald 1996). Our prototype is based on an in-browser automatic DJ mixing framework based on semantic audio technologies (Thalmann et al., 2018), which generates audio feature annotations and uses them to organize and mix together arbitrary audio material.

With this prototype (see Figure 1) users can for example create a timelapse mashup that explores the evolution of one song throughout the history of the band in the form of a diachronic arrangement of all different versions, each playing for the duration of a select number of bars. Alternatively, they can choose to organize the chosen recordings by similarity and overlap them for a significant amount of time in order to create a coherent sound collage reminiscent of John Oswald's. During playback, each fragment is complemented with different visual artifacts found online for the current date. Figure 1 shows a screenshot of the application during a mashup of the song *Me and My Uncle*, currently playing an excerpt of a concert at *The Matrix* while depicting the poster of the original event.

During the experience, users can click on the metadata or artifacts to keep exploring, or they can access the particular recordings they are currently hearing to enjoy them in their entirety. The primary audience of this application are music listeners and fans of the band and the queries and mashups parameters are thus presented as a form of simple presents. However, following the principles outlined above, the platform could be easily adapted to enable digital humanities scholars and musicologists to answer specific questions such as for example finding the most unusual rendition of a particular song, or getting a sense of the audience and atmosphere at a particular venue over time.

Figure 1: The interface of the plunderphonics prototype.

References

Bechhofer, Sean, Kevin Page, and David De Roure. “Hello Cleveland! Linked Data Publication of Live Music Archives,” in Proceedings of WIAMIS, *14th International Workshop on Image and Audio Analysis for Multimedia Interactive Services*, 2013.

Bechhofer, Sean, Kevin Page, David M. Weigl, György Fazekas, and Thomas Wilmering. “Linked Data Publication of Live Music Archives and Analyses.” *Lecture Notes in Computer Science The Semantic Web – ISWC 2017*, 2017, pp. 29–37.

Benson, Michael. *Why the Grateful Dead Matter*. ForeEdge, an Imprint of University Press of New England, 2016.

Oswald, John. “Grayfolded” [CD]. Artifact Music, 1996.

Pampalk, Elias, Andreas Rauber, and Dieter Merkl. "Content-based organization and visualization of music archives." In *Proceedings of the tenth ACM international conference on Multimedia*, pp. 570-579. ACM, 2002.

Thalmann, Florian, Lucas Thompson, and Mark Sandler. "A User-Adaptive Automated DJ Web App with Object-Based Audio and Crowd-Sourced Decision Trees." In *Proceedings of the 4th Web Audio Conference*, Berlin, 2018.

Thalmann, Florian, Thomas Wilmering, and Mark B. Sandler. "Cultural Heritage Documentation and Exploration of Live Music Events with Linked Data." In Proceedings of the 1st International Workshop on Semantic Applications for Audio and Music, ISWC 2018, Pacific Grove, pp. 1-5. ACM, 2018.

Wilmering, Thomas, Florian Thalmann, and Mark B. Sandler. “Grateful Live: Mixing Multiple Recordings of a Dead Performance into an Immersive Experience.” *Audio Engineering Society (AES) 141st Convention*, 2016.