Submitted to *ICCBR'09 Workshop on Reasoning from Experiences on the Web*.

Principle and Praxis in the Experience Web: A Case Study in Social Music

Enric Plaza and Claudio Baccigalupo

IIIA - Artificial Intelligence Research Institute
CSIC - Spanish Council for Scientific Research
Campus UAB, 08193 Bellaterra, Catalonia, Spain.
Email: {enric,claudio}@iiia.csic.es
WWW: http://www.iiia.csic.es

Abstract. The praxis of the users in some particular domain characterizes the Experience Web. In this paper we focus on analyzing *usage of musical objects*, and how the knowledge discovered can be used to design and implement a CBR system. We perform a case study of the Poolcasting CBR system in order to analyze the role of content provided by the praxis of users in the musical domain and the techniques used to acquire the knowledge required by a CBR system in the context of the Experience Web.

1 Introduction

Among the wide diversity of user-contributed content on the web, there is a particular kind of content that has the potential of being put to good use by intelligent systems: *human experiences.* We are now familiar with different forms of content that are provided by the users that reflect not merely an opinion or a belief, but rather express an individual experience: this we may call the Experience Web.

For instance, when a user has experienced a travel with an air carrier company or a stand at a hotel, the comments of that user concerning the air carrier A or the hotel are not merely issues of opinion or belief, they are expressing and recording a concrete and factual experience. That is to say, that the plane was delayed or that the hotel H didn't attend a client's request, they are not merely subjective estimations: (i) they are statements that certain facts occurred and (ii) they are evidence with respect to the likelihood oh these facts being a recurring pattern (A's planes tend to be delayed, H's staff tends to be unfriendly).

Human experiences recorded on the web offer *practical knowledge* concerning a wide variety of real world objects and situations. This practical knowledge is different from theoretical knowledge as that which can be provided by the Semantic Web. For instance, the Semantic Web approach can offer theoretical knowledge about hotels as in the statement "Hotel H is three stars (according to European Standards)", which means that some authority has classified hotel H so because it satisfies certain properties adjudicated to that class. Although this knowledge also provides evidence on the quality and features of hotel H, it is in fact knowledge about the "three star hotel class" rather than about H itself.

Theoretical knowledge is, by definition, about the concept or the class — while practical knowledge is mostly about the object or the instance. What is there about the instance that is not in the class? Well, basically an instance has a concrete *cluster of relationships* with other instances, with its environment (that includes how people use that object or instance). Most of these relations are outside the purview of a theoretical/semantic definition of a concept or a class. Experiences, on the other hand, being concrete, are precisely those *clusters of relationships among instances*.

We have now introduced a core notion: *usage*. Thus human experiences, when expressed, essentially provide a description of *how people have used an object* — and therefore a description of relevant relations of that object with its (physical and conceptual) environment. In previous papers I've emphasized the fact that a large number of experiences in the web are described using text [6, 7]. There are situations, however, where such experiential knowledge are recorded on the web as different forms of data instead of free text; although these situations may seem minority or less general, they may be more amenable to analysis and reuse of those experiences by an automatic process.

In the rest of the paper we will present a case study in the domain of music, where human experiential knowledge is recorded and available as different forms and sources of data, and we will show how this experiential knowledge may be analyzed, interpreted, and reused to automatically to perform a particular task. The task we want to automate is that of play a DJ in a radio channel, as implemented in the Poolcasting system [4]. The task is to convey a selection of songs that are satisfying for a dynamic audience (i.e. a group of individuals) and that play smoothly one after the other (i.e. each song is musically associated with the next and not merely chosen at random). The purpose of the paper is to elucidate how web data can be analyzed as *experiential knowledge* and used in a case-based reasoning process. That is to say, how data can be interpreted as records of actions performed by human beings, and thus represents their *musical praxis*, and how the practical knowledge that can be discovered or inferred from that praxis can be exploited in a system that reasons from people's experiences.

The structure of the paper is as follows. Since Poolcasting has to satisfy two criteria, namely song sequence smoothness and group audience satisfaction, we will show in the next two sections how to acquire experiential knowledge for the criteria of musical smoothness (§2) and audience satisfaction (§3). The paper will close with a discussion section.

2 Social Music Praxis

In order to establish a succession of songs whose order in musically meaningful or appropriate, we need to acquire knowledge about which songs "play well together." Moreover, since this is a matter of degree, we will call *(musical) asso-*

ciation between two songs the likelihood that these two songs "play well one after the other." There are two ways to approach the acquisition of such a *musical association model*, one based on principles, and one based on praxis.

Acquiring a musical association model from *principles* means that we have some theory, some general knowledge, such that when applied to a particular pair of songs yields a degree of association. in the music domain these are called content-based approaches, since they analyze the song's musical or acoustic content. A simple principle could be that songs classified as belonging to the same genre should, in principle, play well together — similarly to the three-star hotel above, we are using information about the class (the genre). We may use any kind of class partition for this purpose, e.g. the songs performed by the same artist, or written by the same composer. The most common way is to represent each song by a collection of acoustic features extracted from the audio signal. For instance, some authors posit the principle that two songs are highly associated when their global timbre quality is similar [1]; this approach then focuses on ways to analyze the spectral shape of songs and ways to assess their similarity. Other approaches use beat and tempo analysis to assess which songs "sound similar" [8].

The other approach is analyzing the praxis of people in situations where they deal with songs that "play well together." One example of this approach would be analyzing the behavior of real DJ's in charge of playing music that flows smoothly over time. Nielsen Broadcasting Data has compiled a large amount of data on music broadcasts in more than 1,600 radio station, but their data base is not available without a fee.

The social web has revolutionized the scope and availability of all forms of content, and specially in the domain of popular music. Different social web platforms that focus on music concerns have compiled user-provided playlists in the order hundred of thousands. A *playlist* is in essence a collection of songs that someone has considered "play well together." Analyzing hundred of thousands of playlists we may discover which songs are more associated with respect to a community of users of a social web platform.

The association model we developed for Poolcasting analyzes this social data to find those songs that appear together in playlists, following two intuitions: (a) that the closer they occur in a playlist, the more associated they are, and (2) the more playlist two songs co-occur, the more associated they are. However, this initial analysis was insufficient, and we needed to take into account what we called a song's *popularity*: the more playlist a song occurs in, the more popular the song is. The reason is that, without taking into account song popularity co-occurrence of pairs of songs in playlists was biased in favor of popular songs. That is to say, we detected mainly situations where one of the co-occurring song was highly popular, but we missed co-occurring pairs of songs where neither of them was popular.

The final measure of *song association*, re-normalizing with respect to relative popularity [2], was applied to 599,565 playlists provided by the social web platform MyStrands.com. Moreover, once song association is estimated we can infer

Destiny's Child	
Poolcasting/ beta=0.5	*Kelly Rowland, City High, Ciara, Fantasia, Christina Milian, Beyonce, Ashanti, Girls Aloud, 3LW, Dru Hill*
MyStrands	*Ciara, Pussycat Dolls, Usher, Beyonce, Nelly, 50 Cent, Mariah Carey, Chris Brown, Gwen Stefani, Eminem*
AllMusic	*Mariah Carey, Jennifer Lopez, Aaliyah, Xscape, Ginuwine, Deborah Cox, Kelly Price, Faith Evans, Brandy, Usher*
Yahoo	*Cruel Story Of Youth, Jessica Simpson, Ryan Cabrera, Ashlee Simpson, Faith Evans, Nick Lachey, Vitaly Romanov, Janet Jackson*
Last.fm	*Beyoncé, Mariah Carey, Jennifer Lopez, Usher, Aaliyah, Rihanna, TLC, Ciara, Ashanti, Christina Aguilera*

Fig. 1. Comparisons of artists associated with Destiny's Child.

association degrees among artists. The results were compared with other models of "musical similarity" among songs and artists, like *All Music Guide* where they are provided by expert editors, *Yahoo! Music* where information come from user feedback, and Last.fm where similarity comes from overall listening habits of users. Although our order-based *association* is asymmetric and is not a measure of *similarity* as the measures provided by these sources, the general results are roughly equivalent in practice. However, our measure did find out more obscure associations (because of the popularity renormalization) than others did not detect. For instance, Figure 1 shows that *Poolcasting* associates Destiny's Child with Kelly Rowland; this association is a good one, because Kelly Rowland is the lead singer of Destiny's Child.

Regardless of the details, the focus of this paper is on the fact that we are analyzing how people *use* their music. Playlists embody some particular instance of the notion of "songs sounding well together," and the social web platform merely provides a conduit where this experiential knowledge, from many users, and about tens of thousands of songs, is expressed and stored. The fact that MyStrands.com is a "social web platform" is relevant in as much as it facilitates that a large number of users contribute their musical experiences. There is no difference in analyzing user's playlists in a personal computer or shared via a website: social web platforms are useful in motivating and facilitating the sharing of experiential knowledge, not necessarily creating that experiential knowledge. Nevertheless, the openness of user-contributed experiential content is very important in practice: (1) the number of songs and artists from different countries and sources (e.g. bootleg concerts, independent bands) is larger than any particular endeavour (like All Music Guide, based on experts) could ever achieve; (2) the responsiveness to include newly created songs is also much higher.

Reasoning from experiences on the web is not only a matter of acquiring and analyzing experiential knowledge. In our musical domain, for instance, playlists are contentless — i.e. they contain references to the songs (and the artists) but not the songs themselves. Thus, to put the Poolcasting system into practice we needed to identify the songs references and the artists names: recover from misspellings, unify denominations and establish a unique (and possibly shared)

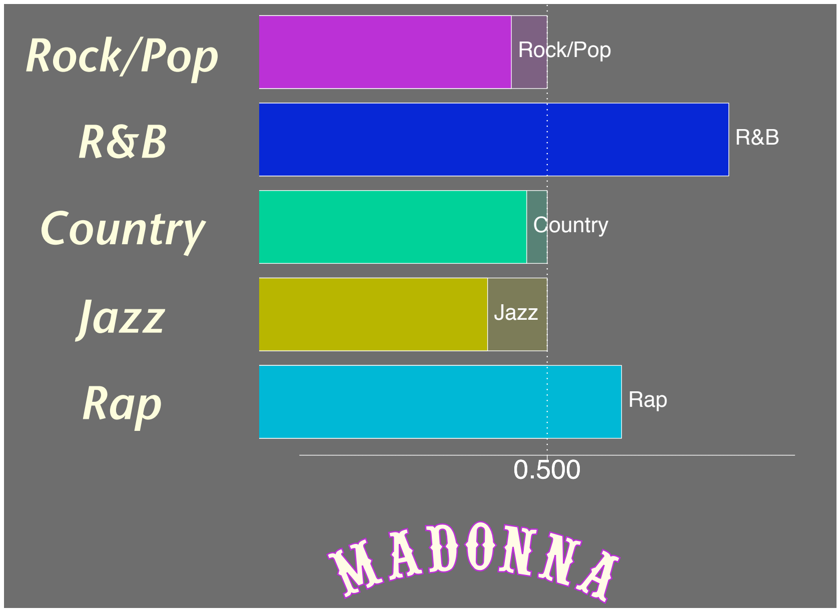

Fig. 2. Genre affinity profile of Madonna.

ID for each object in the domain. This Data Web concern is orthogonal to our approach on the Experience Web, and is currently being addressed by research on the Data Web; in this approach objects like cities or persons are identified by RDF triplets. The "linked data" approach proposed by Tim Bernes-Lee [5] seems more congruous with the Experience Web than the semantic web approach: the Data Web focuses on representing the *clusters of relationships among instances* that we talked about before as the way concrete experiences may be represented.

Analyzing and discovering higher-level relationships from experiences is not technically different from analyzing data, but the fact that the discovered relationship come from data recording practice is what makes a difference. Analyzing how people use and combine songs in playlists we find how songs and artists are associated. Therefore, we can analyze how songs or artists cluster together to form groups. Moreover, songs and artists are already categorized into *genres*, but his application of principles assigns only one genre label to each artist, sacrificing a more nuanced characterization. However, analyzing users' musical praxis we can discover new relations between artists, clusters of artists, and genres.

For instance, we can revise the principle-based categorization of artists and propose that artists have a graded *affinity* to multiple genres [3]. This characterization of artists is closer to reality, since artists do not belong to one genre and are excluded from belonging to any other genre; rather, they have high affinity to some genres (e.g. Madonna has high affinity with Pop and R&B) and low affinity to others (e.g. Madonna has very low affinity with Jazz). Moreover, the

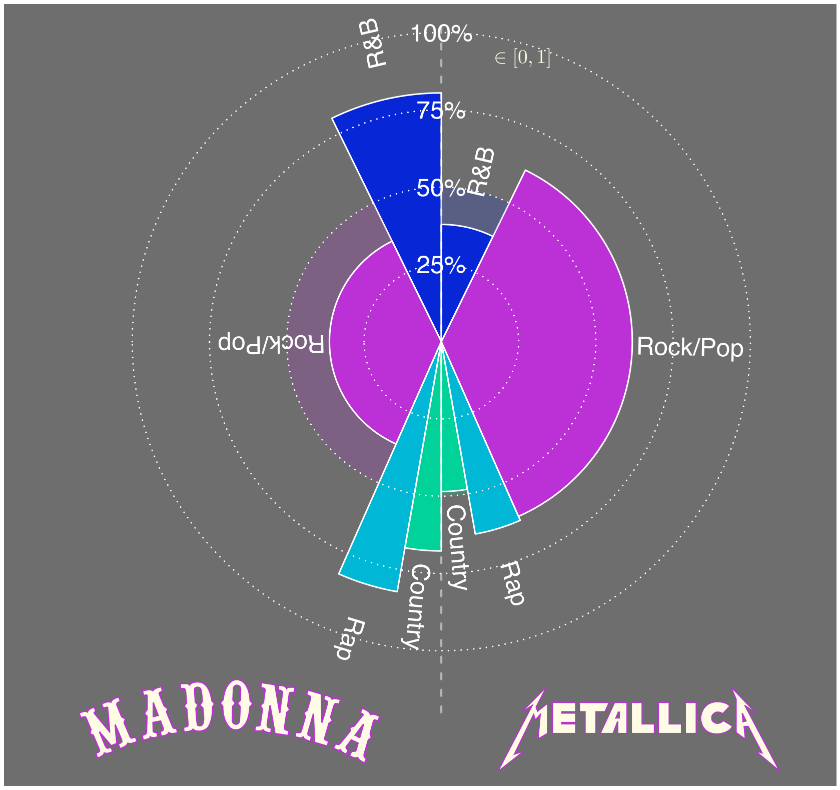

Fig. 3. Genre-centrality comparison of two artists originally labelled as Rock/Pop.

affinity vector of each artist with respect to genres provides a new way to describe musical performers, as shown in Figure 2, where the affinity degree spans from 0 to 1.

Moreover, we can detect which artists are "central" to specific genres — i.e. they are good representatives of that genre. The *genre-centrality* of an artist x to a genre g is the percentage of artists whose genre affinity to g is lower or equal than the genre affinity of x to genre g. For instance, on the Soundtrack genre the most central artists are James Horner, Alan Silvestri and Michael Giacchino, who are famous composers of original movie scores (e.g., James Horners Titanic Original Soundtrack), and not Pop artists who have only sporadically performed famous songs which appeared in movies (e.g., Celine Dions My Heart Will Go On). Moreover, artists can be compared on how central they are to different genres, as shown in Figure 3 where we may compare Madonna and Metallica (both originally classified in the Rock/Pop genre).

The *usage of musical objects* by users is the basis of these analysis, it's the availability of this data that is crucial for the Experience Web. The fact that this data is available in “social web platforms”, or the kinds of data mining techniques used to analyze them, these are secondary issues: the praxis of the users in some particular domain characterizes the Experience Web.

3 Individual Music Listening Praxis

The second criterion for the *Poolcasting* system is to customize the music selection to a dynamic audience —namely the group of users registered at a musical channel. Therefore, the system needs knowledge to estimate how satisfying is selecting a song over another (1) for the individuals in the audience and (2) for the overall satisfaction of the group as such. The overall satisfaction is essentially some sort of average of the individual satisfaction, so we needed to acquire knowledge about which songs and/or artists an individual prefers. For this purpose, we analyzed how individuals used their music libraries on their computers, specifically on iTunes players. The data available on iTunes library database includes which songs are rated higher, which songs had been played frequently, etc.

The strategy is similar to the one in the previous section, but now we are focusing on examining each individual music player as a repository of data about their musical listening praxis. We considered that each library database may be interpreted as an “individual case base” and can thus be used in a CBR system like *Poolcasting* to predict the degree of satisfaction of each user in the audience with respect to a specific song. However, as indicated in the original paper [4], they were not strictly “individual case bases” since they not contained *cases*. The core idea of cases in CBR is very close to that of examples in Machine Learning, maybe for historical reasons; a case may have a representation that is simpler or more complex, but is composed of two *separate* objects: a problem (represented on the *problem space*) and a solution (represented on the *solution space*). Similarity of two problems is defined over the problem space with the purpose of estimating how similar their solutions might be in the solution space.

However, the *Poolcasting* system did not had access to this kind of experiential knowledge represented as cases. Thus, the approach we took was to take a step back, and recall that CBR is also classically defined as reasoning and learning from past experience. Within this interpretation, we did have knowledge of the users’ usage of songs for listening purposes by analyzing the digital music player library data. The knowledge that can be derived is (qualitatively) straightforward: the more often a song has been played, the more star-based ranking a song has, the more songs of an artist the user has, then the more likely is that the user likes that song or that artist.

Moreover, we already have seen in section 2 how to acquire knowledge about song association. The schema in Figure 4 shows how the the *Poolcasting* system combines knowledge about musical association among songs (obtained from playlists contributed by a community of users) and knowledge about audience song preferences (obtained from digital music player library data of the members of the audience). The Retrieve process uses the musical association knowledge to

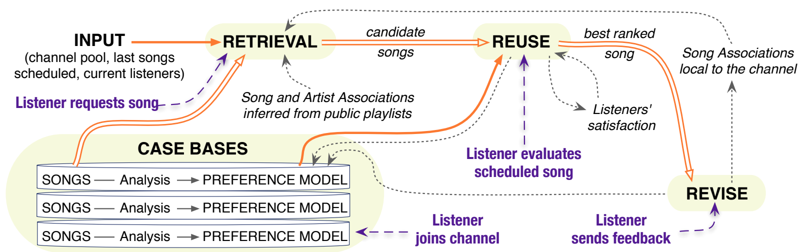

Fig. 4. The CBR schema of the Poolcasting system.

filter, from all possible songs, a small number of songs that are musically associated with the last song being played, while the Reuse process used the audience song preferences knowledge to select the song that will keep the audience (and individual members) satisfied.

4 Discussion

In summary, *Poolcasting* as a CBR system focuses on acquiring and harness knowledge coming from the praxis of users in a domain, analyzing their *usage* of the objects in a domain for specific purposes. In this case study we analyzed how people put together songs in a playlist; they do it because for them these songs (for some unknown reasons or purpose) "sound well together." We also analyzed how individuals use music stored in their digital music players; we interpreted them as repositories of data recording the music listening praxis of each individual.

The Experience Web is therefore characterized by a certain viewpoint on a specific type of content. The content is data representing specific actions, the praxis of individuals in a given domain; the viewpoint is that we interpret those actions, that praxis, as experiences from which new knowledge and insight can be gained and harnessed by developing intelligent systems for achieving specific goals.

The Semantic Web and the Data Web are orthogonal endeavors withe respect to the Experience Web approach. They are in fact required to be able to harness the Experience Web. We have focused on this paper on domains where experiences are directly recorded as data, not free text. Although text-based experiences are qualitatively and quantitatively very important, we would argue that a careful examination of existing non-textual content will uncover areas where the available data can be analyzed and interpreted as experiential content, and be amenable to partake of the Experience Web approach.

References

[1] J. J. Aucouturier and F. Pachet. Music similarity measures: what's the use? In *Proceedings of the 3rd international symposium on music information retrieval (ISMIR02), Paris*, October 2002.

[2] C. Baccigalupo and E. Plaza. Mining music social networks for automating social music services. In *Workshop Notes of the ECML/PKDD 2007 Workshop on Web Mining 2.0*, pages 123–134, 2007.

[3] Claudio Baccigalupo, Justin Donaldson, and Enric Plaza. Uncovering affinity of artists to multiple genres from social behaviour data. In *Proc. 9th Int. Conference on Music Information retrieval ISMIR-08*, pages 275–280, 2008.

[4] Claudio Baccigalupo and Enric Plaza. A case-based song scheduler for group customised radio. In Rosina Weber and Michael M. Richter, editors, *Case-Based Reasoning Research and Development, 7th International Conference on Case-Based Reasoning, ICCBR 2007*, volume 4626 of *Lecture Notes in Artificial Intelligence*, pages 433–448. Springer, 2007.

[5] Tim Berners-Lee. The next web of open, linked data. In *TED2009 Conference*, 2009.

[6] Enric Plaza. Semantics and experience in the future web. In K D Althoff, R Bergmann, M Minor, and A Hanft, editors, *Advances in Case-Based Reasoning: 9th European Conference, ECCBR 2008*, volume 5239 of *Lecture Notes in Artificial Intelligence*, pages 44–58. Springer Verlag, 2008.

[7] Enric Plaza. On reusing other people's experiences. *Künstliche Intelligenz*, 9(1):18–23, 2009.

[8] G. Tzanetakis and P. Cook. Musical genre classification of audio signals. *Speech and Audio Processing, IEEE Transactions on*, 10(5):293–302, 2002.