Chapter 1

Introduction

I ask in all honesty, what would life be?
Without a song or a dance what are we?

ABBA, 1983

1.1   Reusing past experience

Internet is a rich and outspread medium which allows anyone to easily contribute to its content. People share their pictures, videos, songs, texts and, very often, their *personal experiences*, reporting which songs they have played, which friends they have met, which places they have visited, and so on. Sharing personal experience online has become a common activity, which explains the success of Web communities such as Facebook and Twitter offering this kind of service.

Experience data from the Web grants a valuable overview of the social impact of a given content. Browsing the opinions of previous travellers in online forums, for instance, is often more beneficial to decide which hotel to book than the information provided by a guide book. The same occurs with books, videos, music, technical problems, recipes, advices: collecting experiences of others from the Web can offer the knowledge required to perform several tasks.

This dissertation presents poolcasting, an Artificial Intelligence technique that takes advantage of the vast amount of experiential data shared on the Internet to automatically solve a specific task. The task addressed is to customise a sequence of songs for the preferences of a given audience.

There are several situations where people listen to music in groups (discos, radio channels, home-parties) and a professional disc jockey (DJ) is appointed to select the best songs to play for the current audience. Poolcasting is designed to act 'like a good DJ', selecting *automatically* which songs to play from a repository of available music.

While a DJ combines personal expertise and human senses to 'feel' which music is adapt for a given audience, the innovation of poolcasting is to obtain this knowledge from the World Wide Web. Internet makes available a wide repository

of human experiences related to the domain of music: which songs people have played in the past, in which order, how they were rated, commented, forwarded, shared. Poolcasting collects and interprets this kind of human experiences to obtain the knowledge required to deliver smooth and group-customised musical sequences.

The core idea of poolcasting is to solve a task reusing past experiences. This idea characterises a whole family of Artificial Intelligence approaches known as **Case-Based Reasoning** (CBR). CBR systems typically store past experiences as cases, then reuse previous cases to propose solutions for new tasks. A CBR system, for instance, would determine the best cure for a patient by first retrieving similar past problems (patients with the same profiles and symptoms) from a case base, then adapting their solutions (applied cures and effects) to the current case.

Poolcasting represents a reinterpretation of the classical Case-Based Reasoning approach. To solve a task, previous knowledge is reused, but not structured as (problem → solution) pairs. The task to be solved is not to find *one* solution for a given problem, but to iteratively build a good *sequence* of songs that satisfies certain properties in the long run. Poolcasting demonstrates how CBR can make use of the experience of *multiple* users to solve a task that is *social* in nature: to customise content for a group of people.

1.2   A Web of musical data

A focus of this work is on the area of **Web data mining**, describing a technique to collect, analyse, filter and interpret experiential data from the Web to extract valuable knowledge to solve a specific task.

Internet is expanding towards a social medium as more user-generated content is becoming available offering valuable insights about human experiences. This is particularly true in the domain of music.

Music has, on one side, a mathematical nature that has been investigated for centuries and, on the other side, a social nature that has comparatively received less attention. Yet, people commonly influence one another with respect to the music they like and often spend time looking for the 'right' group of people whom to listen music with. On the Web, music lovers can gather in forums to talk about their favourite artists, form virtual fan clubs to show their support, write entries in their blogs and comments in MySpace pages to explain their interests, discover songs that other friends recently played, forward their preferred music to listeners around the world. Music is a universal language, and people of any age, provenience or preference can report their musical experiences of the Web.

One particular type of musical experience data that is widely available on the Internet is playlists. Playlists are sequences of music titles compiled to be played in a specific order. Playlists are useful since they allow people with large music libraries to organise songs in small ordered sequences.

Each playlist can reflect a particular mood or emotion or be shaped by a specific purpose: working, running, cooking, going to sleep. Even ignoring the

motivation that drove someone to put songs in a certain sequence, a playlist indicates that certain songs have been *experienced* together.

Many music-related Web communities invite their members to share personal playlists online, making others aware of songs that are, in their experience, meant to be played together. Millions of playlists are therefore publicly available on the Internet. Poolcasting collects and analyses a large set of these playlists to reveal songs and artists that are correlated according to the people. If two songs or artists co-occur often and closely in multiple playlists, poolcasting understands that those songs and artists share some affinity and it makes sense to play them together. In this way, poolcasting is able to generate sequences of songs that go well together one after the other.

The innovation of this approach is that poolcasting can identify songs that are associated according to the *people*. Other features could be analysed to uncover songs that are similar (e.g., similar lyrics, similar chords, similar acoustic features), but a content-based analysis would not bring a comprehensive overview of songs that people tend to associate. The experiences concealed in playlists, on the other hand, have the power to reveal songs that people associate either for acoustic, social or cultural reasons.

1.3   Listening to music in a group

Another focus of this thesis is related to the research in the area of **social choice**. The purpose of poolcasting is to customise music for a *group of listeners* (an audience), and this activity implies looking for an acceptable compromise among the musical preferences of the entire audience.

The most appreciated disc jockeys are those that can build 'good' musical sequences for every type of audience, making all the participants move to the dance floor during an event, and not just a few people. Similarly, the goal of poolcasting is to deliver music that can possibly satisfy the entire group.

If a group were made of people with identical musical tastes, it would be easy to find songs that *everyone* likes. This scenario, though, is quite rare: most groups show different individual preferences for different songs, and what someone likes, someone else dislikes.

Under these conditions, the approach of poolcasting is to build a musical sequence that fairly satisfies all the listeners *in the long run*. At times, people may be exposed to songs they do not particularly like but, after a certain while, everyone will be satisfied by the overall music played.

The way in which poolcasting tackles this social issue is by introducing a new preference aggregation method that assigns different importance to different members according to how much satisfied they are towards the music played so far. This satisfaction-weighted aggregation method is designed to select, at each moment, songs that are most liked by the least satisfied listeners so that, in the long term, a balance can be reached in the entire audience.

1.4   A social radio experience

While the first part of the thesis describes the poolcasting technique, the second part introduces an application developed on top of this technique. The application, called Poolcasting Web radio, provides an innovative online radio service that streams music channels on the Internet, where songs are selected in real time according to the preferences of the connected listeners.

In the real world, people are used to listen to music in groups. In bars, cars, elevators, open offices, people share musical experiences and implicitly accept a social contract by which their musical interests only partially influence the music played, which is intended to satisfy the group as a whole. This kind of social experience does not occur on the Internet.

When listening to music on the Web, people typically choose between online radios or digital music services. Online radios are equivalent to terrestrial radios; they just stream over the net rather than over the air and have a much larger number of channels. They are not social, since connecting to a channel does not bring any information about who else is listening, nor allows like-minded listeners to know each other. Moreover, online radios are not personalised: they broadcast either random or pre-programmed musical sequences that are not influenced by the listeners.

The alternative to online radios is provided by digital music services such as Last.fm and Pandora which stream personalised music channels. These services compile a music profile of each member and select songs according to the preferences of each individual. Again, they lack of a social component since streams can only be listened by one person at the time and cannot be shared.

The fact that the Web is becoming more and more a *social* medium and that online radios and digital music services are only targeted to *individuals* is the motivation that led to the development of Poolcasting Web radio, which offers an online *social radio* experience.

Poolcasting Web radio is not meant to help people share music *files*, but to share *musical experiences*, allowing displaced persons to listen to the same music at the same time, to exchange their interests, to contribute together to the sequence of music played and to be able to discover new music that others like.

1.5   Structure of the thesis

Each chapter of the thesis touches a particular research field; for this reason the state of the art related to each area is reported separately at the beginning of each chapter.

Chapter 2 is dedicated to the experience Web and explains how Internet can be mined for data related to musical experiences in the form of playlists, and how playlists can be analysed to learn which songs and artists are more associated according to the 'wisdom of the crowd'. Previous work about co-occurrence analysis and musical associations is also presented.

Chapter 3 explains how to generate individual music profiles from listening behaviour data. Most digital music players store data about individual listening history (which songs a person has played and rated) and many users agree to share these data on the Web to make others aware of their musical experience. This chapter first introduces previous work about user modelling and then describes a technique to analyse listening behaviour data to infer a model of the musical preferences of each listener.

Chapter 4 illustrates the core of the dissertation: the poolcasting technique. Given a group of listeners and a repository of songs, poolcasting employs the techniques of Chap. 2 and 3 to collect the knowledge required to determine a sequence of songs adapted to a given audience. Poolcasting is characterised by an iterated CBR process that first *retrieves* good candidate songs to be added to the sequence, next *reuses* the acquired knowledge to identify the best candidate, then *revises* the knowledge to customise the solution for the current problem. The chapter reviews state of the art about Case-Based Reasoning and group-adaptive systems and also introduces the satisfaction-weighted aggregation method employed by poolcasting to achieve fairness in the long run.

Chapter 5 first reviews previous work about group-adaptive music systems and Internet radios and then describes Poolcasting Web Radio, the innovative application developed to provide an online social radio service. Each channel of the radio is driven by a process that selects in real time which songs to broadcast according to the preferences of the listeners in the same channel at each moment. The radio offers the chance to share music and listening experiences to listeners located around the world.

Chapter 6 reports the evaluation of the presented work, showing the degree in which musical sequences can be customised for a given audience while maintaining a certain musical continuity from song to song. Different scenarios are compared, measuring the performance of poolcasting with groups of different size and degree of homogeneity in terms of musical interests.

Chapter 7 summarises the contributions of this research work and suggests ways to extend poolcasting to domains other than music, in order to deliver customised sequences of movies, news items or TV shows to groups of people.