Semantic descriptors to help the hunt for music

January 4 2006

You like a certain song and want to hear other tracks like it, but don't know how to find them? Ending the needle-in-a-haystack problem of searching for music on the Internet or even in your own hard drive is a new audio-based music information retrieval system.

Currently under development by the SIMAC project, it is a major leap forward in the application of semantics to audio content, allowing songs to be described not just by artist, title and genre but by their actual musical properties such as rhythm, timbre, harmony, structure and instrumentation. This allows comparisons between songs to be made and listeners to find little-known tracks that suit their tastes but may otherwise go unnoticed.

"The music world is highly commercial and only the works of the biggest artists are really well known and widely promoted," notes SIMAC project manager Xavier Serra at Barcelona's Pompeu Fabra University. "Something like 10 per cent of music accounts for 90 per cent of music sales, while the remaining 90 per cent accounts for just 10 per cent of sales – this system could therefore herald a revolution for little-known music and artists."

Technologically, the project has made significant progress toward bridging the so-called 'semantic gap' in audio content. Existing classification systems, such as those used to generate play lists on popular PC media players and MP3 devices, are based on low-level description techniques using text data about the artist, the track, album

and genre. Users can find other tracks by the same artist, from the same album or within the same genre but there is no guarantee that the songs will be anything but remotely alike. Alternatively, websites and online stores selling music often provide recommendations to users based on their personal preferences and the past purchases they and other customers have made.

"The 'since you bought this artist, you might also want to buy this one, as other customers with a similar profile did' method is not that effective, however, because there may be similar songs out there but if they haven't been bought that much they won't appear as a recommendation," Serra says.

The SIMAC approach, which incorporates machine learning, signal processing and musical knowledge as well as text retrieval, overcomes these problems by using a technique to describe music by its actual properties and characteristics. The musical facets of songs are analysed and the tracks are automatically tagged by an annotator. This then allows them to be organised based on the similarities between them within a music surfer programme.

"The technique represents a major advancement over the existing methods used by audio software," Serra notes. "It improves the way users can organise, navigate and visualise audio files and how they can interact with music on their audio player, PC or the Internet."

Another component of the system is a music recommender for users to obtain recommendations that really interest them about new or old songs from online stores. Called FoaFing the Music, it uses not only the musical characteristics of songs to recommend similar ones but also the users' profile, their past purchasing history and what has been written about the songs in website news and reviews. It is based on the Friend of a Friend (FoaF) concept that draws on information from thousands of

machine-readable Web pages via RSS feeds.

The prototype has drawn "positive feedback" from trial users, according to Serra, and commercial interest in the SIMAC project, which ends in March, is high.

"The system offers evident advantages to users in the way they can find and interact with music, and big benefits to artists, producers and the music content industry as a whole," he notes. "Lesser known artists and small production companies who don't have big promotional budgets should benefit in particular because it will increase their visibility and the accessibility of their content."

The system could also be a boon to the software and consumer electronics industries, Serra predicts, noting that there is extensive interest in employing the annotation and music surfing techniques in Internet music sites, PC software and portable devices.

Project partner Philips is currently developing an MP3 player incorporating audio analysis components and one SIMAC component has already been licensed to mSoft, a US company that will use it to search library music to find authorless tracks and sound effects that cannot be categorised using traditional methods.

"Though we were originally planning to set up a spin-off company to exploit the project results, we're now looking to collaborate with other firms to introduce our technology into existing products and others that are still in development," Serra says.

Source: IST Results