



© 1997–2004, Millennium Mathematics Project, University of Cambridge.

Permission is granted to print and copy this page on paper for non-commercial use. For other uses, including electronic redistribution, please contact us.

16/01/2006

News

Lost in music



These are great times if you're a music lover. The World Wide Web opens the door to a near infinite volume of music, and new technology allows you to store a sizeable chunk of it on your hard drive or carry it around with you in a little box. The problem is that choice is time-consuming. You could spend half a lifetime looking for new artists and bands on the Web, and the other half organising what you have downloaded. But luckily, this is about to change.



A very old hat

Scientists working on the SIMAC project (Semantic Interaction with Music Audio Contents) have developed a new audio-based music retrieval system which can compare and classify songs and recommend new music based on the user's taste. Existing systems, like those that put together playlists on computers or MP3 players, choose songs according to information such as name of track, artist, album and genre. Online shops and website recommend music based on the customer's previous choices and those of others that have bought the same track or album. But none of this information is inherently musical everything is based on text that someone, somewhere has previously entered. Songs that are deemed similar may seem worlds apart to another user, and little-known artists that have never been classified can slip through the net.

The new SIMAC system is the first to consider data coming directly from the music itself. "The technique represents a major advancement over the existing methods used by audio software," says SIMAC project manager Xavier Serra, "it improves the way users can organise, navigate and visualise audio files and how

Lost in music

they can interact with music on their audio player, PC or the Internet."

To achieve its aim, the system has to "listen and learn": it has to make sense of the immensely complicated audio signal coming from the music, combine what it learns with other information, and then classify the song. Getting a machine to listen to music and extract information about features like beat, rhythm and harmony is an intense mathematical problem. The system has to search the audio signal for patterns and periodicities, using a technique called *audio signal processing*. With the help of some heavy mathematical machinery (for example a *Fourier transform*) the numerical information which represents a piece of music can be searched for patterns that reveal frequencies and amplitudes. Additionally, stochastic methods can be used to search out structure in songs.

But what these methods typically come up with is some pretty raw technical data about the music not at all what ordinary users like you and me want from a music retrieval system. What we would like the system to do is to classify songs and recognise which ones are similar, so that it can put together playlists and recommend new music from the Web.

The SIMAC prototype bridges the "semantic gap" between technical data and the end user by cleverly combining signal processing techniques, machine learning and musical knowledge. Its "semantic descriptors" analyse music according to facets like rhythm, harmony and timbre. Algorithms search each facet for features that help to characterise a song's genre and mood: the use of percussion instruments, for example, or the typical intro–verse–chorus–verse–chorus–outro structure of pop songs.

When it comes to comparing songs, the system allows a lot of user input, since similarity is in the eye of the beholder and even varies over time: today you might be in the mood for Jazz, while tomorrow you might want only happy songs, irrespective of genre. The system contains a visual interface that lets you pin–point exactly what "similar" means to you at this precise moment, and then off it goes to come up with a selection.

The recommendation component combines data that comes directly from the audio with information from the user's profile and information about artists, which it finds on the Web. It's called *FoaFing the music* because it's based on the *Friend of a Friend* (FOAF) Internet project, which creates a web of machine readable pages that describe people, their interests and the links between them. The team believe that this is the first time that both musical and "cultural" criteria are used to make recommendations.

Project partner *Phillips* is currently developing an MP3 player that will use SIMAC techniques, and a US company has licensed a SIMAC component for a system that searches libraries of sound effects that can't be easily identified. But the prototype may not only improve the lives of music listeners. According to Serra, it could also turn the music industry on its head: "The music world is highly commercial and only the works of the biggest artists are really well known and widely promoted," he says, "Something like 10 per cent of music accounts for 90 per cent of music sales, while the remaining 90 per cent account for just 10 per cent of sales this system could therefore herald a revolution for little–known music and artists."

Marianne Freiberger



Plus is part of the family of activities in the Millennium Mathematics Project, which also includes the NRICH and MOTIVATE sites.