

Music, CompuMaestro -- like Radiohead, please...

September 15 2008



Figure 1: The data flow and system modules for automatic accompaniment

Songwriting: A Schematic. Credit: USC Viterbi School of Engineering

It's an archetypal exchange in musical performance. A vocalist stands poised to perform. The guitarist alongside is ready to add depth and harmony to the melody. The guitarist doesn't know the song but "hum a few bars and I'll fake it," she tells the singer. "Could you do it in the style of Radiohead?" asks the vocalist. "No problem," the versatile guitar player says.

Now a software system created by two University of Southern California researchers can do the same -- not just create an appropriate accompaniment, but do so in the style of any chosen artist, or even the

particular style used in select pieces by the artist. The system can potentially run on an ordinary PC.

Elaine Chew, an accomplished pianist who is a professor at the USC Viterbi School Department of Industrial and Systems Engineering, and graduate student Ching-Hua Chuan began working on their ASSA (Automatic Style Specific Accompaniment) system two years ago.

Chuan is a guitarist who has played with a number of rock bands in her native Taiwan; she received a PhD in computer science from the Viterbi School in 2008. Her prizewinning dissertation presented much of the ASSA research.

The aim, according to a paper presented a year ago at the International Joint Workshop on Computational Creativity in London, was as straightforward as it is ambitious: "... we describe an automatic style specific accompaniment system that makes songwriting accessible to both experts and novices. ... [T]he system should be able to identify the features important to the style specified by the user, [enabling the user to] ask for harmonization similar to some particular songs."

The ASSA system meets the challenge, according to both subjective and rigorous statistical tests.

In the London presentation, titled "A Hybrid System for Automatic Generation of Style-Specific Accompaniment," Chuan and Chew laid out the basics of the system, and tested it on Radiohead songs. They trained the system on three Radiohead songs, and generated chord progressions for the fourth; the original (i.e. Radiohead's own) accompaniment served as "ground truth" — that is, the test for the aptness of the accompaniment.

Thus, the measure of success was not a subjective perception of whether

the accompaniment worked, but rather, based on the observation that it was very close to the band's own accompaniment.

Since then, they have gone on to a more sophisticated system, which they will present in September at the International Conference on Music Information Retrieval. Using as raw material songs from Green Day's "Dookie" and "American Idiots," Keane's "Hope and Fears," and Radiohead's "Pablo Honey" and "Hail to the Thief," they altered various elements of the ASSA modeling technique to observe how the output changes.

The takeoff point of the Chew-Chuan analysis is one on which much recent music theoretic work is based — the work of 19th century German theorist Hugo Riemann, who rethought the old ideas of harmony, creating new ways of representing rules that make some, but not all, rhythmic successions of overlaid notes sound musical. Riemann's representation has only recently been applied to rock music.

The ASSA analysis looks at the tree of possible accompanying chord sequences, and analyzes which branches are followed, treating the unfolding string of possibilities as a Markov series – that is, a temporally ordered series of states (in this case, chords) in which chord determines its successor.

The resulting system thus has two elements. One is analysis of complete, accompanied music samples for what (neo-Riemannian) harmonic progressions they tend to follow. The second is taking these derived stylistic rules and using them to generate accompaniment -- using a Markov series -- for a user-provided melody.

The original London paper lays out the steps in the ASSA attack on the problem. "The system first determines the chord tones in the melody. [One ASSA] module applies machine learning techniques to choose the

chord tones from the input melody, based on the [chosen stylistic] example pieces. The system uses 17 attributes to represent the melody."

Then, chords are prescribed at checkpoints in the melody where [the system finds] the harmony unambiguous. "Using these checkpoints as anchors, we use neo-Riemannian transformations to build chord progressions between checkpoints. Finally, we use Markov chains to generate the final chord progression."

The initial study took four songs by Radiohead: "High and Dry," "Fake Plastic Trees," "Airbag," and "Creep" and applied the rules derived from analyzing three of the songs to generate an accompaniment for a fourth. That is, the bare melody of song four was the starting point, to which the stylistic rules derived from the other three songs were applied to create a new, ASSA accompanied version.

"Creep" and "High and Dry" were both used as tests. Analyzing, the "overall correct rate" of the chord tone choices made by ASSA was 82 percent (for a 54-note sample) in "Creep," and 70.5 percent on a 61-note sample for "High and Dry."

The tests used to evaluate the effectiveness of its accompaniment were at least partially subjective, based on listeners' opinions in Turing tests, in which they had to guess if the accompaniment was the original or machine-generated.

ASSA is not the first attempt to create a robotic accompaniment. I-Ring (by Hong-Ru Lee and Jyh-Shing Roger Jang of National Tsing Hua University) and MySong (by Ian Simon, Dan Morris, and Sumit Basu at Microsoft Research) generate accompaniment based on training sets of 150 and 298 songs respectively -- songs from various genres are treated homogeneously as a whole for training the systems. David Temperley and Daniel Sleator created another system, the rule-based Harmonic

Analyzer in Melisma.

Only ASSA and MySong aim to emulate style, and in MySong, style is classified into two broad modes: "happy" and "jazz." The ASSA approach treats style at a much more individual and specific level – as the property of pieces from a particular period in a band's output, or as features unique to one individual piece.

In a follow-on paper, Chuan and Chew propose a more rigorous and objective set of statistical measures, and use these tests to compare ASSA against the Harmonic Analyzer and a naïve system that requires only that accompanying chords contain at least one of the melody notes that it harmonizes.

This paper, "Evaluating and Visualizing Effectiveness of Style Emulation in Musical Accompaniment," to be presented at the Ninth International Conference on Music Information Retrieval at Drexel University in Philadelphia on September 15, proposes six metrics, three based on straight one-to-one comparisons of the machine created accompaniment with the original — "correct rate, same chords, chords-in-grid" — and three statistical tests measuring musical closeness.

In an experimental test on music by Green Day, Keane, and Radiohead, Chuan and Chew trained their ASSA system altering various parameters, including the number of similar songs included in the rule finding sample, chord similarities between training songs and target melody, and others.

Again, the idea was to first train the system using a few similar songs from a single artist, then to take the melody of another song from the artist, have the system produce an accompaniment, and then compare that accompaniment to the "ground truth" original.

This time the comparison was done with statistics, rather than human listeners. They find that the statistics support the human choices in the earlier study. Among the results: ASSA selected more "right" chords, and choosing songs with more common chords with the target created better results.

And, intriguingly: For most machine learning systems, "it is often the case that more training data guarantees better results. But for ASSA goals, more training songs often do not improve the output quality." In style emulation, having more data appears to dilute the style specificity of the output. Chuan and Chew find that training the system on any one song in the same album resulted in better accompaniment than training it on all songs in the album.

Chuan and Chew plan to continue the work. Among the possibilities is building an end-to-end interactive prototype that will take user humming input, create the accompaniment, and render it with the appropriate instruments.

As the system is now, "some parts of the code are in Matlab and other parts use WEKA," a machine learning tool, said Chuan. "Next we will implement the system as a standalone program either in Java or C++."

The processing demands are not heavy: "A PC is definitely sufficient for this program," she said.

To hear samples of the accompaniments created, go to: www-scf.usc.edu/~chinghuc/Site/Samples.html

The two papers referred to can be downloaded in pdf form:

C. H. Chuan and E. Chew, "A Hybrid System for Automatic Generation of Style-Specific Accompaniment," 4th International Joint Workshop on

Computational Creativity, London, UK, June 17-19, 2007.
www-rcf.usc.edu/~echew/papers/CC2007/chuanchew_cc07.pdf

C. H. Chuan and E. Chew, "Evaluating and Visualizing Effectiveness of Style Emulation in Musical Accompaniment," the 9th International Conference on Music Information Retrieval, Philadelphia, PA, September 14-18, 2008.

ismir2008.ismir.net/papers/ISMIR2008_179.pdf

Source: University of Southern California

Citation: Music, CompuMaestro -- like Radiohead, please... (2008, September 15) retrieved 23 April 2024 from <https://phys.org/news/2008-09-music-compumaestro-radiohead.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.