Representation and discovery of feature set patterns in music

Mathieu Bergeron and Darrell Conklin

Abstract. In this paper we describe and apply a method for discovering patterns within music corpora. Previous approaches to the problem have represented melodies in terms of single abstract attribute sequences, and therefore efficient repeated substring algorithms could be applied for the detection of recurrent patterns. This paper presents a new approach, based on sequential pattern mining and knowledge representation methods for music, that permits data mining of music event sequences in which events are represented as multiple feature sets. A feature set represents a conjunction of features held by an event, and can refer to any number of rhythmic and melodic attributes. A depth-first maximal frequent sequence pattern mining algorithm was implemented and applied to melodies of French singer-songwriter Georges Brassens. Results illustrate the feasibility of extending music pattern discovery methods with a more expressive pattern representation.

Melodic Contour Extraction for Indian Classical Vocal Music

Ashutosh Bapat, Vishweshwara Rao, and Preeti Rao

Abstract. The problem of pitch tracking of the singing voice in the presence of Indian percussive interference, specifically the tabla, is considered. To overcome the problems due to this particular type of interference, a pitch tracker is used that applies dynamic programming (DP) based smoothing on pitch estimates obtained from a spectral-domain pitch detection algorithm (PDA) that uses harmonic matching. Experiments on real and simulated signals show the superiority of the spectral domain PDA over a correlation domain PDA in terms of pitch detection accuracy and suitability of the PDA output for post-processing. A new smoothing cost function is proposed and evaluated. The paper formulates general rules guiding the choice of cost functions participating in the DP based post-processing for this particular problem.

A cooperative approach to style-oriented music composition

D. Espi, P. J. Ponce de Leon, C. Perez-Sancho, D. Rizo, J. M. Inesta, F. Moreno-Seco, and A. Pertusa

Abstract. Evolutionary methods have been largely used in algorithmic music composition due to their ability to explore an immense space of possibilities. The main problem of genetic related composition algorithms has always been the implementation of the selection process. In this work, a pattern recognition-based system helped by a number of music analysis rules is designed for that task. The fitness value provided by this kind of supervisor (the music critic) models the affect for a certain music genre after a training phase. The early stages of this work have been encouraging since they have responded to the a priori expectations and more work has to be carried out in the future to explore the creative capabilities of the proposed system.

A Complexity-based Approach to Melody Track Identification in MIDI Files

Søren Tjagvad Madsen and Gerhard Widmer

Abstract. In this paper, we will examine the importance of music complexity as a factor for melody recognition in multi-voiced popular music. The assumption is that the melody (or lead instrument) will contain the largest amount of information – that it will be the least redundant voice. Measures of melodic complexity calculated from pitch and timing information are proposed. We test the different complexity measures and different prediction strategies, and evaluate them on the task of predicting which track of a MIDI file contains the main melody. Filtering out melody tracks can be useful when searching large databases for similar songs. 108 melody track annotated pop songs were included in the experiment.

A Genetic Rule-based Expressive Performance Model for Jazz Saxophone

Rafael Ramirez and Amaury Hazan

Abstract. In this paper, we describe an evolutionary approach to inducing a generative model of expressive music performance for Jazz saxophone. We begin with a collection of audio recordings of real Jazz saxophone performances from which we extract a symbolic representation of the musician's expressive performance. We then apply an evolutionary algorithm to the symbolic representation in order to obtain computational models for different aspects of expressive performance. Finally, we use these models to automatically synthesize performances with the expressiveness that characterizes the music generated by a professional saxophonist. This work extends our previous work on inducing rules using genetic algorithms for understanding expressive music performance.

An Experiment on Evolutionary Essentic Sound

Soh Igarashi and Koichi Furukawa

Abstract. This paper describes an experiment on evolutionarily generated essentic sound as a new proposal for composing evolutionary music. Essentic sound originally proposed by Dr. Manfred Clynes is the sound that expresses several emotions human beings typically have. This work utilized genetic algorithm for synthesizing essentic sound and tried to generate such sound material as expresses an emotion gradually rising from some chaotic sound. The algorithmic details are given and its aesthetic intention is discussed. The significance and potentiality of this work is also discussed. Some examples of the synthesized sound can be obtained through the author's web page.

Using Multiple-Part Learning and a Genetic Algorithm to Compose Emotion-Inducing Tunes

Roberto Legaspi, Yuya Hashimoto, Koichi Moriyama, Satoshi Kurihara, Masayuki Numao

Abstract. This work aims to computationally specify a system's music compositional intelligence that tightly couples with the listener's affective perceptions. The system initially learns first-order rules from multiple-part data to model a listener's emotional impressions of music using musical score fragments that were hand-labeled by the listener according to a semantic differential scale that uses six affective adjective pairs. A genetic algorithm, whose fitness function is based on the acquired rules and follows specific music theory, is then used to generate variants of frame and chord structures. Lastly, the system integrates melody to the GA-obtained structures. Empirical results show that the system is 80.6% accurate at the average in classifying the affective labels of musical structures and is able to automatically generate musical pieces that stimulate four kinds of impressions, namely, favorable-unfavorable, bright-dark, happy-sad, and heartrending-not heartrending, which the listener can distinctly perceive.

The Critical Damped Oscillator Fitness Function in Music Creativity Problems

Tzimeas Dimitrios, Mangina Eleni

Abstract. This paper discusses the contribution of Genetic Algorithms (GAs) to music creativity and focuses on the design of the fitness function highlighting main aesthetic problems that occur from different algorithmic strategies. It presents a simple music problem as a base for

three sets of experiments illustrating the performance of different choices of fitness function. It introduces a novel genetic algorithm technique where the fitness function is assigned dynamically based on the simulation of the critical damped oscillator. Given a valid data representation the fitness function operates as a force of a critical damped oscillator keeping successfully the chromosome close to a predefined set of targets.

A Probabilistic Context-Free Grammar for Melodic Reduction

Édouard Gilbert and Darrell Conklin

Abstract. This article presents a method used to find tree structures in musical scores using a probabilistic grammar for melodic reduction. A parsing algorithm is used to find the optimal parse of a piece with respect to the grammar. The method is applied to parse phrases from Bach chorale melodies. The statistical model of music defined by the grammar is also used to evaluate the entropy of the studied pieces and thus to estimate a possible information compression figure for scores.

Optimal Parameter Set Acquisition for exGTTM

Yoshinori Oka, Masatoshi Hamanaka, Keiji Hirata, and Satoshi Tojo

Abstract. Although GTTM (Generative Theory of Tonal Music) has been considered to be one of the most promising theories in musical structure retrieval, its rules are often ambiguous and conflictive. We have implemented an automatic time-span tree analyzer (ATTA) , where the weight or the priority of each rule could be changed with adjustable parameters. However, because those parameters were only controlled manually, it took so much time for us to find a set of optimal parameters. In this paper, we propose a parameter tuning system with Genetic Algorithm (GA), as a simultaneous multi-points searching. As the crossover strategy of GA, we employ UNDX (Unimodal Normal Distribution Crossover) to optimize the performance. As a result, the search by GA outperformed other naïve search methods.

Preliminaries for Transformational Analysis in OpenMusic

Yun-Kang Ahn, Carlos Agon, and Moreno Andreatta

Abstract. This paper deals with the implementation of transformational analysis in OpenMusic. This theory offers new perspectives for musical theory and analysis, introducing graphs and networks. OpenMusic, based on a visual programmation paradigm, is a powerful score editing software that can also be used for musical analysis. Its graphic environment might provide graph drawing, computer-aided decision and motivic detection that would allow us to interest in the segmentation problem and to analyze scores according to this new musical theory.

Sung Note Segmentation for a Query-by-Humming System

Pradeep Kumar, Manohar Joshi, Hariharan, S. Dutta-Roy, Preeti Rao

Abstract. Retrieval performance in query-by-humming (QBH) systems depends crucially on the accurate note segmentation and labeling of user queries. To facilitate note segmentation, querying is often restricted to the easily detected syllable “ta”, which is not necessarily the syllable most preferred by users. In this work, new acoustic features based on the signal energy distribution as obtained from the singing perception and production points of view are investigated. Performance evaluations on a manually labeled database of syllabic humming show that a specific mid-band energy combined with a biphasic detection function achieves high correct detection and low false alarm rates on the sonorant consonant syllables /da/, /la/ and /na/. The resulting onset detector is incorporated in the signal-processing front-end of an available QBH system (hitherto constrained to ta-syllable queries only). QBH retrieval performance results are reported on a large dataset of user queries.

The Non-Trivial Machine in Digital Audio

Arun Chandra

Abstract. An implementation of von Foerster's non-trivial machine in software, for the generation of simultaneous waveforms that modify their own structure based on the current structure of their neighbors. This research creates a contemporary version of counterpoint, but one that takes advantage of the computer's ability to avoid notes. Thus, instead of "note-against-note" behavior (as was the case with 16th- and 17th-century counterpoint), this research explores "wave-against-wave" behavior. The construction of each waveform allows for the linking of square, triangular and curved elements in any combination. Thus, a variety of timbres is the structural output of the waveform. Each element in each waveform has from six to twelve parameters, and on every iteration each parameter is linearly changing. So, the frequency, amplitude and timbre of the waveform is gradually shifting. Sudden change can occur if neighboring waveforms reach defined thresholds.

Score-Guided Music Audio Source Separation

Christopher Raphael

Abstract. We present a technique for music audio source separation using a time-aligned symbolic musical score. Separation is achieved by inverting a time-frequency audio representation using only those time-frequency points primarily dedicated to a particular instrument or group. Our experiments concern the problem of separating a violin soloist from a full orchestra. We first treat the problem with real audio data in which the contributions from soloist and orchestra are known. In this case we label time-frequency points using a classifier, learned from labeled training data, whose accuracy can be measured. We then extend these results to incorporate realistic constraints on the labeling. This latter method is tested on data from a commercial compact disc.

Clustering Streaming Music via the Temporal Similarity of Timbre

Jacob Merrell, Dan Ventura, and Bryan Morse

Abstract. We consider the problem of measuring the similarity of streaming music content and present a method for modeling, on the fly, the temporal progression of a song's timbre. Using a minimum distance classification scheme, we give an approach to classifying streaming music sources and present performance results for auto-associative song identification and for content-based clustering of streaming music. We discuss possible extensions to the approach and possible uses for such a system.