

The Melody Triangle - Pattern and Predictability in Music

Henrik Ekeus (1), Samer Abdallah (1),
Mark D. Plumbley, Peter W. McOwan
(1) Centre for Digital Music
Queen Mary University of London

ABSTRACT

The Melody Triangle is a musical pattern generating system. It has two interfaces; one is a traditional screen based interface, the other a multi-user interactive installation. In both cases, the Melody Triangle allows its users to interactively explore patterns of predictability in music. It makes use of statistical models developed as part of Information Dynamics of Music (IDyOM) project[ref], which seeks to model patterns of expectation and surprise in the perception of music.

We outline the Information Dynamics model and how it forms the basis of the Melody Triangle. We discuss both uses of the system, the multi-user installation where collaboration in a performative setting provides a playful yet informative way to explore expectation and surprise in music, and the screen based interface where the Melody Triangle becomes compositional tool. Finally we outline a study where participants used the screen-based interface under various experimental conditions to allow us to determine the relationship between the Information Dynamics models and musical preference. We found that...

Keywords

Information dynamics, Markov chains, Collaborative performance, Aleatoric composition

1. INTRODUCTION

Music generally involves patterns in time. Composers commonly, consciously or not, play with his or her audience's expectations by setting up patterns that seem more or less predictable, and thus manipulate expectations and surprise in the listener[ref]. The research into Information Dynamics explores several different kinds of predictability in musical patterns, how human listeners might perceive these, and how they shape or affect the listening experience.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NIME12, Somewhere

Copyright 2008 Copyright remains with the author(s).

2. INFORMATION DYNAMICS AND THE TRIANGLE

(some background on IDyOM and Markov chains)

The Information Dynamics model operates on discrete symbols, only at the output stage is any symbol mapped to a particular note. Each stream of symbols is at any one time defined by a transition matrix. A transition matrix defines the probabilistic distribution for the symbol following the current one. In fig.1, on the left we see two transition matrices, the upper one having no uncertainty and therefore outlining a periodic pattern. The lower one containing considerable unpredictability but is nonetheless not completely without perceivable structure, it is of a higher entropy. The current symbol is along the bottom, and in this case there are twelve possibilities. In the upper matrix in fig. we can see for example that symbol 4 must follow symbol 3, and that symbol 10 must follow symbol 4, and so on. Hundreds of transition matrixes are generated, and they are then placed in a 3d statistical space based on 3 information measures calculated from the matrix, these are redundancy, entropy rate, and predictive-information rate [see [cite]]. In fig.1 on the right, we see a representation of these matrixes distributed; each one of these points corresponds to a transition matrix. Entropy rate is the average uncertainty for the next symbol as we go through the sequence. A looping sequence has 0 entropy, a sequence that is difficult to predict has high entropy rate. Entropy rate is an average of surprisingness over time.

Redundancy tells us the difference in uncertainty before we look at the context (the fixed point distribution) and the uncertainty after we look at context. For instance a matrix with high redundancy, such as one that represents a long periodic sequence, would have high uncertainty before we look at the context but as soon as we look at the previous symbol, the uncertainty drops to zero because we now know what is coming next.

Predictive information rate tell us the average reduction in uncertainty upon perceiving a symbol; a system with high predictive information rate means that each symbol tells you more about the next one. If we imagine a purely periodic sequence, each symbol tells you nothing about the next one, that we didn't already know as we already know how the pattern is going. Similarly with a seemingly uncorrelated sequence, seeing the next symbol does not tell us anymore because they are completely independent anyway; there is no pattern. There is a subset of transition matrixes that have high predictive information rate, and it is neither the periodic ones, nor the completely un-corellated ones. Rather




Figure 1: Screen Shot from the Information Dynamics engine - the current and next transition matrixes are on the left. The upper one has no uncertainty and thus represents a periodic pattern. The lower one contains unpredictability but nonetheless is not completely without perceivable structure. On the right we see the population of transition matrixes distributed along three axes of redundancy, entropy rate and predictive -information rate. Note how the distribution makes triangle-like plane floating in 3d space.

they tend to yield output that have certain characteristic patterns, however a listener can't necessarily know when they occur. However a certain sequence of symbols might tell us about which one of the characteristics patterns will show up next. Each symbols tell a us little bit about the future but nothing about the infinite future, we only learn about that as time goes on; there is continual building of prediction.

When we look at the distribution of matrixes generated by a random sampling method in this 3d space of entropy rate, redundancy and predictive information rate, it forms an arch shape that is fairly thin, and it thus becomes a reasonable approximation to pretend that it is just a sheet in two dimensions(see fig.1). It is this triangular sheetfig.2 that is then mapped either to the screen, or in the case of the interactive installation, physical space. Each corner corresponding to three different extremes of predictability/unpredictability, which could be loosely characterised as periodicity, noise and repetition.

3. USER INTERFACES

The Melody Triangle engine, developed in Prolog and MatLab, can be controlled by OSC messages and thus any number of interfaces could be developed to for it. Currently two different interfaces exist; a standard screen based interface where a user moves tokens with a mouse in and around a triangle on screen, and a multi-user interactive installation where a Kinect camera tracks individuals in a space and maps their positions in the space to the triangle.

3.1 The Multi-User Installation

the statistical properties of this melody is based on where in the physical room the participant is standing, as this is mapped to a statistical space (see below). By exploring the

physical space the participants thus explore the predictability of the melodic and rhythmical patterns, based on a simple model of how might guess the next musical event given the previous one. ...

When multiple people are in the space, they can cooperate to create musical polyphonic textures. For example, one person could lay down a predictable repeating bass line by keeping themselves to the periodicity/repetition side of the room, while a companion can generate a freer melodic line by being nearer the 'noise' part of the space.

A Kinect camera was used to tack individuals in the space, an application developed in OpenFrameworks would send the individuals positions and a bounding box values (for gesture recognition), to an application running the Information Dynamics Engine [Matlab/Prolog]. This would map the input to a series of symbols that would go over OSC to Max/MSP, where these would be interpreted as MIDI, mapped to channels and the data passed on to Logic for output using software instruments.

3.1.1 Tracking and Control

It uses the OpenNI libraries' API and high level middleware for tracking with Kinect. This provided reliable blob tracking of humanoid forms in 2d space. By then triangulating this to the Kinect's depth map it became possible to get reliable coordinate of visitors positions in the space.

This system was then further extended to detect gestures. By detecting the bounding box of the 2d blobs of individuals in the space, and then normalising these based on the distance of the depth map it became possible to work out if an individual had a an arm stretched out or if they were crouching.

With this it was possible to define a series of gestures for controlling the system without the use of any controllers.




Figure 2: The Melody Triangle - the triangle's axis corresponds to

Thus for instance by sticking out one's left arm quickly, the melody doubles in tempo. By pulling one's left arm in at the same time as sticking the right arm out the melody would shift onto the offbeat. Sending out both arms would change instrument.

3.2 The Screen Based Interface

[screen shot] On the screen is a triangle and a round token.

With the mouse you can click and drag the red token and move it around the screen. When the red token is dragged into the triangle, the system will start generating a sequence of piano notes. The pattern of notes depends on where in the triangle the token is

4. MUSICAL PREFERENCE AND INFORMATION DYNAMICS STUDY

The study was divided in to 5 subtasks. The axes of the triangle would be randomly rearranged prior for each participant.

this first task, which will last [4/3] minutes, we simply ask you to move the token where ever in the triangle you wish,. This allowed the participant to get use to the environment get use to the interface and get a sense of how position of tokens changes a melody.

In the following tasks a background melody is playing and the participant are asked to find a second melody that 'works well' with the background melody. In each of these tasks the background melody has different statistical properties. In the first it, In the second the background melody ... in the third... And finally in the fourth case the melody is in the middle of the triangle, that is it....

5. RESULTS

X participants took part in the study (mean median age).
(Prior musical experience?)

5.1 Stats of results

6. OBSERVATION/DISCUSSION

7. ACKNOWLEDGMENTS

This work was partly supported by two EPSRC grants: GR/S82213/01 and EP/E045235/1.




Figure 3: ?On the left we see the depth map as seen by the Kinect, and the bounding box outlines the blobs detected by OpenNI. On the right is an eagle-eye view of the position of individuals tracked in the space. These are the positions that are mapped to the statistical space of the information dynamics engine.