Title
Outside the Framework of Thinkable Thought: The Modern Orchestration Project

Author
Gattegno, Eliot

Publication Date
2010

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Outside the Framework of Thinkable Thought: The Modern Orchestration Project

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Musical Arts

in

Contemporary Music Performance

by

Eliot Aron Gattegno

Committee in charge:

Professor Steven Schick, Chair
Professor Jordan Crandall
Professor Charlie Oates
Professor Katharina Rosenberger
Professor Rand Steiger

2010

The Dissertation of Eliot Aron Gattegno is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2010

TABLE OF CONTENTS

VITA

2005	Bachelor of Music, New England Conservatory of Music, Boston
2007	Master of Music, New England Conservatory of Music, Boston
2010	Doctor of Musical Arts, University of California, San Diego

ABSTRACT OF THE DISSERTATION

Outside the Framework of Thinkable Thought:
The Modern Orchestration Project

by

Eliot Aron Gattegno

Doctor of Musical Arts

University of California, San Diego 2010

Professor Steven Schick, Chair

In today's world of too much information, context -not content- is king. This proposal is for the development of an unparalleled sonic analysis tool that converts audio files into musical score notation and a Web site (API) to collect manage and preserve information about the musical sounds analyzed, as well as music scores, videos, and articles on music. By combining currently available software with custom created programs in a novel way, this project revolutionizes the field of music by empowering general and expert users with the ability to put information in a context never before imaginable. Applying the Wikipedia model, user created data will be

available through a Web site to users at all times for no cost. The data will be searchable through innumerable variations of custom filters -like those of a travel Web site- and employ a newly conceived poetic search engine that has question-answering capabilities currently unavailable. Eventually, this database will become the largest source of information (including sounds, scores, articles, and recordings) on musical instruments and scholarship in the world. This project defines innovation by radically changing the way researchers search and obtain information, therefore, changing the way they think about information, and changing the way they create with information. In summary, the project leverages what computers do best – organize, store and analyze data – permitting humans to spend more time doing what they do best – art.

Chapter 1

In January 2010, this venture commenced as a joint research project amongst doctoral students at the University of California, San Diego and Harvard University. A project summary based on the past ten months of research and software trials is outlined over the following pages.

Currently, there are many efforts to develop centralized databases on the Internet. Those that are most significant to this project are JSTOR and ARTstor. JSTOR, is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive of over one thousand academic journals and other scholarly content. JSTOR uses information technology and tools to increase productivity and facilitate new forms of scholarship. ARTstor's mission is to use digital technology to enhance scholarship, teaching, and learning in the arts and associated fields. ARTstor consists of a repository of hundreds of thousands of digital images and related data; tools to actively use those images; and a restricted-usage environment that seeks to balance the rights of content providers with the needs and interests of content users. What both of these databases have in common is they have become the primary aggregator for guidance in their respective areas of inquiry. This project plans to use their model to achieve similar success in the field of music. Also of interest is a computational knowledge engine currently under development, Wolfram|Alpha. Wolfram|Alpha's long-term goal is to make all systematic knowledge immediately computable and accessible to everyone. The project aims to collect and curate all objective data;

implement every known model, method, and algorithm; and make it possible to compute whatever can be computed about anything. Like Wolfram|Alpha, this project takes advantage of all currently existing technologies and models, improves on their success, thus, offering and anticipating the needs of it's users, sometimes before they even knew they had them.

This project will be developed in three phases: Phase one's result will be the creation of first-of-it's-kind sonic analysis software that converts audio files into musical score notation. This phase will take three months. Outlined below are the technical aspects associated with this phase.

The project will be created entirely with free, open-source software (FOSS), using well established software languages and tools which already have large communities of users and developers. The core of the project will be written in Python, because of its flexibility, ease of authoring, and the vast amount of third-party programming materials already available in the language. A variety of FOSS Python libraries, including NumPy (Numeric Python), SciPy (Scientific Python), Gamera (a document analysis library which implements an efficient and flexible KDTree - a way of searching high-dimensional spaces), Matplotlib (for creating charts and plots, especially of audio analyses) and Abjad (a library for formalized musical score control) will all be utilized to both analyze, search and display the data collected. Professor Michael Casey's (of the Graduate Program in Digital Musics at Dartmouth College) command-line tool, fftExtract (which itself wraps the FFTW library, "the fastest Fourier transform in the West"), will be employed as a method of extracting salient audio features from the corpus of recorded instrumental events. GNU Lilipond, a FOSS application for music engraving, will be used for creating all graphical representations of music notation. Furthermore, a Linux-based web-server running Apache, MySQL and PHP (including the FOSS Zend Framework) will be employed to create the public face of the project. Finally, all of the source code will be made free and available to the public, especially the wider academic and artistic community, for further study, development and refinement.

Phase two's result will be a Web site (API - Application Programming Interface) that will manage the database of sounds, scores and written materials. This will take place over a period of three months. Outlined below is information on how the Web-based public Front-end of the site will serve as well as how the poetic search engine will have the ability to perform a multi-dimensional search of the orchestrational corpus to be created in phase three.

Given the complexity of the project, the enormity of the amount of data it aims to both collect and produce, and the arcane nature of much of the sonic analyses involved, a concise and comprehensible interface for non-technical users is an absolute necessity. To that end, the project intends to make a publicly accessible web-interface, both as a simple means for perusing the collected information, as a tool for performing high-level searches and compositional activities, and as a social mechanism for commenting on the extant material and adding new. As many of the instrumental techniques to be documented are unstable, or vary wildly from instrument to instrument (for example, between different brands of woodwinds), the project gives the wider musical community the ability to add further examples to the orchestrational corpus, to better present an increasingly inclusive view of the landscape. The project also provides various web-based affordances for non-technical users, so they can search the database intuitively, using the same sorts of techniques as more technical users, but without necessarily needing to understand the intricacies of those techniques. They should be able to, for example, click on the various parts of a woodwind instrument fingering diagram to describe a range of possible fingerings to look for, or they could highlight a portion of a spectrogram and then adjust a slider to find other sounds with similar spectral regions within a certain "distance".

The relative complexity of searching for anything in a database is intricately linked to the number of dimensions involved, and significantly increases with that number. Consider searching for restaurants near a street address in any contemporary geographic search service such as Google Maps, or MapQuest: the question of how "near" a particular restaurant is to a given address is only one of the physical distance (the Euclidean distance, really) of the restaurant's address to the target address, based on just two dimensions - latitude and longitude. However, how can one get an understanding of "nearness" in a "musical space," where there are considerably more than two dimensions, in fact often many dozens, and where the number of dimensions one can consider useful varies wildly according to the intentions of the researcher? The problem of searching in these extraordinarily high-dimensional spaces is of considerable concern to the broader computer science community, as modern relational databases are unable to efficiently query information by "distance" in dozens or even hundreds of dimensions. Given the inherent ambiguities of the musical search domain, the project is designed to be as flexible, efficient and clear as possible in the structuring of it's search tools - to provide as many avenues and combinations of avenues for searching as desired, to ensure that the searches operate quickly, and to allow the researcher (or simply, user) to approach their task with whatever level of complexity they desire.

Phase three uses the sonic analysis tool and Web site (API) created in phase one and two to gather essential information on a critical mass of musical instruments to produce a corpus sufficient for sovereign use. This final phase will take sixth months. Outlined below is a comprehensive look at how both sonic and idiomatic data will be analyzed.

This projects aim is to bring as much of contemporary sonic analysis to bear on the audio we collect as possible. Additionally, we propose to analyze not only the audio, but also the idiomatic content of the performance techniques themselves; that is, the patterns of fingerings and breath that make the sounds possible. The sonic techniques to be used include, but are by no means limited to Mel-frequency-scale cepstral coefficients, a psycho-acoustically grounded description of timbre (timbre being a general musical concept of "sound color"), chromagrams, a method for describing the overall harmony of a sound, as well as lower-level descriptors such as log power, log harmonicity, spectral centroid, spectral flatness, spectral slope, and so forth. All of these are tested and proven techniques, in use today in many academic and commercial applications. Idiomatic information - that is, instrumental fingerings -

can often be represented as simple text strings, and then compared and analyzed as text. The Levenshtein edit-distance algorithm, often used to perform fuzzy-string-matching (or to suggest alternate spellings when spell-checking a document), can be employed to detect similar sets of performance information. The project hopes that in accumulating a critical mass of information, and then analyzing that information in as many ways as possible, never-before-seen musical possibilities will be discovered: that the clusters and patterns that appear in our research will lead the way to new sounds. In order to have a way to teach and perfect these new sounds, the project staff is already researching ways to allow the automatic generation of pedagogical materials. The project intends to use the data collected, and the tools developed to manage and search that data, to create new pedagogical materials for performers and composers, both as a means to stimulate interest in some of the more obscure sound worlds the project wants to showcase, and to make comprehensible those materials through the more traditional garb of standard pedagogical practices. To that end, the project will to programmatically generate instrumental etudes, using the data we have collected as well as a battery of pattern-creating techniques. These etudes, in different levels of complexity to match the skill of the performer - from grade school through professional, would be typeset so as to appear identical to standard notation practices, and would therefore allow performers to learn new techniques in the way they are already accustomed to.

For initial support, the project intends to approach the Andrew W. Mellon Foundation, as well as the National Endowment for the Humanities for a Digital

Humanities Start-Up Grant. In addition to the foundation and grant funding, the project has the potential to support itself by offering premium service subscriptions to institutions and libraries. Lastly, the project staff plans to approach instrument manufacturing companies for sponsorship. These companies will have access to research specifically performed on their instruments for use to further develop their instrument manufacturing techniques.

In order for the project staff to evaluate the project effectively, several measures have been taken to ensure success. The first is the construction of an internationally recognized advisory committee to oversee the project's development and growth. The project staff intends to receive regular feed back from the board, and have budgeted time within the previously proposed limits to incorporate their feedback. Using Google Analytics, the project staff will be able to observe how users are interacting with the Web site. Also, the interactive and communal nature of the site offers ample opportunity for users to ask questions (of the staff and other users), provide feedback, and make suggestions for improvement for future versions of the project.