CHARACTERIZING COMPOSERS USING JSYMBOLIC2 FEATURES

Cory McKay
Marianopolis College
cory.mckay@mail.mcgill.ca

Julie Cumming
McGill University
julie.cumming@mcgill.ca

Ichiro Fujinaga
McGill University
ichiro.fujinaga@mcgill.ca

ABSTRACT

Statistical features extracted from large collections of symbolic music can be of important musicological value. This utility is explored here via several experiments involving Renaissance composers: training machine learning models to identify the composer of a piece; statistically analyzing feature values in order to learn what empirically differentiates compositional styles; and using machine learning to investigate composer attribution validity. This work also demonstrates the potential of the jSymbolic2 software and the features it extracts.

1. INTRODUCTION

Composer attribution is much more than a toy problem in early music studies, as there are many pieces whose composer is unknown or disputed. The combination of automatically extracted features and machine learning provides significant potential for resolving such debates.

In addition, statistical feature analyses can provide musicologically important insights on compositional styles, as they allow one to empirically study all available music by a given composer based on a broad spectrum of characteristics. This contrasts with traditional manual research, where time limitations necessitate a focus on just a few pieces and a limited range of characteristics. A large automatically extracted feature catalogue may also arguably permit more consistent and "objective" research than can be achieved by human experts alone, even the best of whom cannot avoid at least a little bias. At a minimum, such features provide a fresh perspective, by allowing music to be considered in ways that might be unintuitive but still highly useful.

Automatically extracted features also facilitate studies with a breadth and scope that would take several lifetimes of often tedious manual work to accomplish. These can be used both to investigate the empirical validity of existing theories and to perform purely exploratory studies.

2. JSYMBOLIC2

jSymbolic2 is open-source Java software that extracts features from digital symbolic music files. In addition to the types of automatic classification and statistical analysis tasks described here, these features can also be used for additional purposes, such as content-based search, clustering, visualization, etc.

172 unique features can be extracted by jSymbolic2, for a total of 1230 feature values (some features are multi-dimensional vectors). These features are associated with a wide range of musical characteristics, including pitch statistics, melodic intervals, vertical intervals, chords, rhythm, instrumentation, texture and dynamics. Although there are several other excellent symbolic music analysis platforms (e.g. MIDI Toolbox, music21 and Humdrum), none of them include a feature catalogue anywhere near as large or broad as that of jSymbolic2.

jSymbolic2 is also intended to serve as an easily extensible platform researchers can use to implement and extract their own bespoke features. A plug-in architecture is used to facilitate this extensibility.

jSymbolic2 is a dramatically expanded version of the original jSymbolic [2], and can be downloaded for free from http://jmir.sourceforge.net. We would like to thank SSHRC and the FRQSC for their generous funding for the development of this software and related work.

3. COMPOSER ATTRIBUTION EXPERIMENTS

Our first set of experiments involved using supervised learning to classify music by its composer. We constructed the new "RenComp7" dataset, which is a combination of: the Palestrina music collected by John Miller; the Victoria music collected by Jon Wild and Andie Sigler; and the Josquin, de la Rue, Ockeghem, Busnoys and Martini music from [3]. This resulted in 1584 MIDI files.

726 of the 1230 jSymbolic2 features (chosen to avoid bias based on source) were then extracted from the RenComp7 pieces, and 10-fold cross-validation experiments were performed using Weka's SMO support vector machine implementation (with default hyper-parameters).

Several such experiments were performed, one involving classifying amongst all seven RenComp7 composers, and nine more focusing on certain composer subsets of particular musicological interest. Table 1 details the impressive classification accuracies achieved.

The excellent work of Brinkman et al. [1] provides the best available published context for these results. The authors used 53 features to classify between 6 composers (J. S. Bach and five Renaissance composers), and obtained success rates of roughly 63% on average.

RenComp7 Composers	Average CV Accuracy (%)
All 7	92.7
Ockeghem/Busnoys/Martini	87.2
Ockeghem/Busnoys	84.4
Ockeghem/Martini	94.6
Busnoys/Martini	93.8
Josquin/Ockeghem	93.9
Josquin/Busnoys	96.0
Josquin/Martini	88.2
Josquin/de la Rue	85.4
Victoria/Palestrina	99.9

Table 1: Composer identification accuracies averaged across cross-validation folds.

YES: Less music for more than 4 voices
YES: More 3-voice music
YES: More triple meter
SAME: Less stepwise motion
SAME: More notes at the bottom of the range
SAME: More chords (or simultaneities) without a third
SAME: More varied rhythmic note values
OPPOSITE: More large leaps (larger than a 5th)
OPPOSITE: More dissonance

Figure 1: Empirical validation of expert predictions as to musical characteristics that would be more evident in Ockeghem's music than Josquin's. "YES" means the expectations were empirically correct, "SAME" indicates no statistically significant difference between the two and "OPPOSITE" means the expected characteristic was more associated with Josquin than Ockeghem.

Rodin Certainty Level	Percent Classified as Josquin
Level 3	48.6
Level 4	17.2
Level 5	14.0
Level 6	5.5

Table 2: Percentage of pieces associated with each of Rodin's Josquin attribution certainty levels classified using jSymbolic2 features as in fact being by Josquin.

4. FEATURE ANALYSIS RESULTS

We next used the jSymbolic2 features to empirically examine the stylistic differences between pairs of composers. In particular, we compared features extracted from all available music by (the somewhat different) Josquin and Ockeghem, and (the quite similar) Josquin and de la Rue.

We began by looking at how well the data supported the expectations of two Renaissance music experts, Julie Cumming and Peter Schubert, as to what characteristics would differentiate the styles of Josquin and Ockeghem. The results, shown in Figure 1, demonstrate how some of their predictions were indeed correct, but others were not. This underlines the general need for these kinds of empirical investigations into the validity of a wide range of musicological and theoretical beliefs and assumptions.

Weka was then used to apply seven statistical feature analysis techniques to highlight the features that most effectively distinguish the composers in each pair, and the results were combined into a ranked feature list that revealed new musicological insights. It turns out that a combination of rhythmic characteristics are particularly important in distinguishing Josquin from Ockeghem and, furthermore, Ockeghem tends to have more vertical sixths and diminished triads, as well as longer melodic arcs. With respect to Josquin and de la Rue, Josquin tends to have: more vertical unisons and thirds; fewer vertical fourths and octaves; and more melodic octaves.

5. JOSQUIN ATTRIBUTION CERTAINTY

There is musicological debate as to whether certain pieces associated with Josquin were truly composed by him. Jesse Rodin has helpfully broken Josquin's music into six attribution certainty categories, where Level 1 is the most secure and Level 6 the least [3]. We investigated this empirically by training an SMO model on jSymbolic2 features extracted from the music of 21 Renaissance composers taken from the JRP [3] and the two most secure Josquin levels (i.e. levels 1 and 2). The music in the remaining four Josquin levels was then classified as being either by or not by Josquin using this model.

Table 2 shows that the more insecure the category, the less likely a given piece was to be classified as being by Josquin. Although these results must be taken with a very large grain of salt (e.g., they would be more meaningful if a much wider range of non-Josquin music had been used to train the classifier), they do demonstrate some partial empirical support for Rodin's categorization.

6. CONCLUSIONS AND FUTURE WORK

This work demonstrates the significant research potential of the jSymbolic2 features when applied to musicological research on composers. Future work will involve expanding this study by experimenting with music by more composers and more music per composer. Additional areas of musicological interest beyond composers will also be studied; comparing madrigals with motets is an immediate goal. We will also apply the jSymbolic2 features to other kinds of music, including non-Western and popular musics. The jSymbolic2 feature catalogue is currently being expanded with these goals in mind.

7. REFERENCES

[1] A. Brinkman, D. Shanahan and C. Sapp, "Musical stylometry, machine learning and attribution studies: A semi-supervised approach to the works of Josquin," *Proc. of the Biennial Int. Conf. on Music Perception and Cognition,* pp. 91–97, 2016.

[2] C. McKay, "Automatic music classification with jMIR," Ph.D. diss., Schulich School of Music, McGill Univ., Montreal, Canada, 2010.

[3] J. Rodin and C. Sapp, Josquin Research Project [Online]. Available http://http://josquin.ccarh.org.