Modeling Rhythmic Complexity in a Corpus of Polyrhythm Examples from Europe and America, 1900-1950

Ève Poudrier[1] and Daniel Shanahan[2]

[1] *School of Music, University of British Columbia, Canada*
[2] *School of Music, Ohio State University, United States*

[1]eve.poudrier@ubc.ca, [2]daniel.shanahan@gmail.com

Abstract

Rhythmic complexity, as represented by *polyrhythm* (the superposition of two or more contrasting rhythms, meters or speeds), is often identified as one of the central features of twentieth-century Western art music. This project uses computational analysis to explore the development of polyrhythm in a corpus of 719 examples extracted from 450 works by composers from Europe and North America from 1900 to 1950 (Suter, 1980). The current study aims to develop complexity metrics to examine the use of polyrhythm in this period, test competing claims about its development, and explore the cognitive processing of complex rhythms. Corpus examples and associated metadata were processed to be analyzed using the *Humdrum Toolkit* (Huron, 1995). Exploratory analysis was conducted using a stratified sample dataset ($N = 80$) that includes four randomly selected examples for each composer ($N = 20$). Correlational analysis using global complexity measures (average nPVI, event density, entropy) showed that differences in entropy and variability between rhythmic groups within each excerpt could be predictive of genre, and that the entropy of the composite rhythms seems to decline in the 1930s and 1940s, before increasing again in the 1950s. These results are taken as starting points for future avenues of research on the interaction of rhythmic groups, variability, and complexity in the use of polyrhythm in Western art music of the twentieth century.

Introduction

While the cognitive mechanisms that support entrainment to periodic auditory signals have become increasingly well-defined (e.g., Large & Jones, 1999), little research has focused on how musical agents (which may be identified as belonging to one or more of the three categories traditionally described as listeners, composers, and performers), in various cultures, actually perceive and aesthetically evaluate complex musical rhythms. Such rhythmic complexity is often identified as one of the central features of twentieth-century Western art music. In particular, the use of *polyrhythm*, that is, the superposition of two or more contrasting rhythms, meters or speeds, seems to have increased significantly both in prevalence and scale. Before the turn of the twentieth century, apart from a few isolated cases (e.g., Mozart's simultaneous use of three different orchestras, each playing in a different meter in a scene from *Don Giovanni*, 1787), polyrhythms were generally used locally or as a means to create special accompaniment textures (e.g., Chopin's Etude No. 1 from *Trois Nouvelles Études*, 1840).

Several competing ideas have been offered to explain this stylistic development. Composers' growing interest in experimentation, including the representation of multiple temporal perspectives simultaneously, has been interpreted as stemming from a desire to express more intense human experiences brought about by urbanization and social unrest. Alternatively, the resulting musical structures have been characterized as mechanistic, reflecting the rapid development of industrialization and technology. The bulk of evidence presented in support of these claims is mostly anecdotal in nature, often taking a composer's statement as a point of departure and mapping it loosely onto socio-cultural trends or historical events. Yet another approach, common in music-theoretical research, has been to use close analysis to generate a set of observations, which are then formalized based on some guiding principles borrowed from a variety of disciplines such as acoustics, mathematics, linguistics, or phenomenology. Common to these approaches is the reliance on close reading, subjective interpretation, and generalization from a small number of examples believed to be representative (e.g., Krebs, 1999; Poudrier, 2009; Roeder, 1994), without the possibility to test the validity of the claims, that is, to measure the probability that the findings are not due to chance or researchers' bias.

The current study is part of a larger project that seeks to develop a framework for the computational analysis of polyrhythm, and to explore the development of polyrhythmic techniques in a corpus of 719 musical examples extracted from 450 works by composers from Europe and North America in the first half of the twentieth century, when the use of these techniques has been noted to increase significantly. The primary research objectives of this project are to: (1) identify trends in the development of polyrhythm over the focus period within the context of musical production (composer's national origin, date of composition, and location of first public performance as well as performers involved); (2) identify structural features and associated musical parameters, such as timbre, register, dynamics, and performance manner (e.g., staccato vs. legato); and (3) explore psychometrics associated with polyrhythm perception. The current study focuses on the first objective and aims to develop complexity metrics to examine the use of polyrhythm in this corpus, test competing claims about its development, and explore the cognitive processing of complex rhythms.

The Suter (1980) Corpus

The source corpus was compiled in print format by Louis-Marc Suter (1980). Suter's study aimed to identify trends in stylistic development as well as to classify the specific structures and effects resulting from the use of polyrhythmic techniques in the works of twenty representative composers in the period of 1900 to 1950, from Leoš Janáček (1854-1928) to Benjamin Britten (1913-1976). While Suter's study provides

evidence for an increase in the use of polyrhythm, both in terms of prevalence and complexity, the large number of works surveyed and the methodology he adopted (close reading of a sub-set of examples) place important limitations on the reliability and specificity of his findings, and thus, many of his claims must be regarded as speculative.

So far, computational analysis of temporal structures has been limited to rhythmic patterns that are integrated in a single metric hierarchy, including the disruption of an established meter (Huron & Ommen, 2006; Temperley, 1999; Volk, 2008). Suter's sampling of polyrhythm examples in 913 representative works presents a unique opportunity to expand the scope of rhythm research, and to test the validity of claims about the development of polyrhythm in this specific style and period. The encoding of this corpus will also facilitate the design of ecologically valid stimuli for listening experiments aimed at exploring the psychophysical mechanisms and cognitive constraints at work in the processing of complex rhythmic structures, and could provide a basis for comparative studies of polyrhythm production and perception in different cultures (e.g., such as those found in African diaspora music).

Methods

Sampling of source materials. Suter's original corpus is comprised of 913 works by twenty different composers from Europe and North America (see Figure 1), ranging from Leoš Janáček (1854-1928) to Benjamin Britten (1913-1976). To allow for generalization across the various styles and genres represented in the focus period, Suter selected these twenty composers based primarily on the timing of their output (i.e., majority of works composed between 1900 and 1950) and even distribution in terms of composers' age during the focus period.
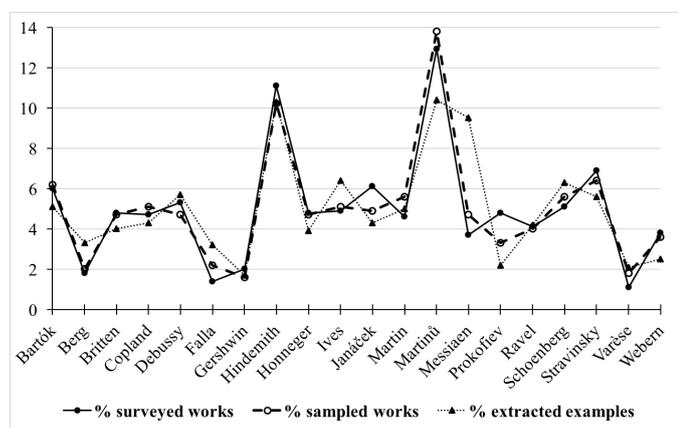

Figure 1. Composer percentage representation in the Suter (1980) Corpus based on surveyed works, sampled works, and extracted examples. A total of 913 of works were surveyed; 719 examples were extracted from 450 of these works. Composers' average representation within type is 5% (SD = 2.6 to 2.9), with works and examples by Martinů (10.4 to 13.8%) and Hindemith (10.2 to 11.1%) being the most represented, and those by Varèse (1.1 to 2.1%), Gershwin (1.6 to 2%), and Falla (1.4 to 3.2%) being the least represented.

Among the group of available composers that satisfied these basic criteria, specific composers were selected based on their recognized influence on the music of the period ("individual character") and the diversity of genres

represented in their output ("stylistic diversity") rather than on the prevalence of polyrhythm in their works. There was no attempt at equal representation in terms of national origin. To allow for characterization, surveyed works were then selected to provide a representative sample of the full range of a given composer's creative output, both in terms of style and genre, rather than proportional representation across composers or genres.

These selection criteria resulted in an unbalanced dataset in terms of national origin (e.g., there are four French composers, but no Italian composer), genre (larger proportion of orchestral and instrumental works in comparison to vocal or stage works), and composer representation (number of works/examples per composer, as shown in Figure 1). Examples are also quite varied in terms of length, ranging from 1 to 44 notated bars ($M = 6$, $SD = 4$), while source work full units (e.g., movement or scene from which an example was extracted) range from 9 to 1,044 notated bars ($M = 231$, $SD = 192$), with extracted examples representing an average of 5.5% of the source work full units ($SD = 9.2$). The current corpus is comprised of the examples for which score excerpts were included in Suter's dissertation; these were extracted from 450 works, with 151 of these works being represented by two or more extracted examples.

Encoding of examples. A *reserve dataset* (RDS) was created using a stratified sample of the full corpus (designated as FDS) to allow for both exploratory analysis and hypothesis testing. The RDS is comprised of four randomly selected examples for each composer ($N = 80$), the remaining 639 examples forming the *testing dataset* (TDS). Examples from Suter (1980) that featured non-adjacent segments were split into separate examples resulting in a total number of 719 examples for the current corpus. The encoding of the corpus involved transcribing the score excerpts using Sibelius so that they could be converted in kern representation using Humdrum. For works that featured texted parts, all musical data were transcribed except for the text, and to facilitate score reading, transposing instruments were transcribed to sounded pitch (in C). Transcribed examples were then proofread and edited manually as needed, including correction of translation errors. To date, all of the RDS examples have been encoded, and 101 examples from the TDS are still in process, including eight examples that could not be transcribed due to software limitations and some examples for which a published score has not been located. It has not been determined yet whether these examples can be manually encoded, and whether they will be excluded from the corpus or will be retained for the purpose of metadata analysis.

Metadata gathering. For the purpose of exploring correlations between structural features and aspects of musical production, metadata related to the twenty represented composers as well as to each example's source work and first public performance were collected and linked to the primary data (see Table 1). The primary sources for metadata collection consisted of in-print as well as online authoritative scholarly sources such as scholarly editions of a composer's complete works, Oxford Music Online, and biographies as well as other easily accessible resources such as composer-dedicated websites, IMSLP, CD booklets from Naxos Music Library, and Wikipedia. Metadata collected from non-

scholarly sources were subjected to a data validation process requiring confirmation by at least one additional source. Additional metadata pertaining to each example's temporal attributes were collected using the most easily accessible printed edition of each work. Recorded duration was collected using randomly selected recordings from the Naxos Music Library (one for each example). To date, metadata gathering for the full dataset (FDS) is at 93.6% average completion, with most fields having reached above 90% completion, except for premiere organization (46.7%), premiere venue (67.2%), and premiere ensemble (84.3%).

Table 1. Corpus metadata based on composers, works, premieres, and examples. Sub-type categories are shown in parentheses.

Composers	• Birth/death place (city, country) • Birth/death year • Nationality
Works	• Genre (orchestral, instrumental, vocal, stage) • Sub-genre (symphonic, concerto, programmatic, keyboard, chamber, song, choral, opera, ballet, theatrical) • Composition year (first, last) • First publication year
Premieres	• First public performance year • Location (organization, venue, city, country) • Performers (conductor, ensemble, soloists)
Examples	• Source work full unit[a] • Notated length • Recorded duration • Time signature • Tempo (expression, notated beat, metronomic rate)

[a]For multi-movement works (and stage works), this corresponds to the specific movement (or scene) where the excerpt is located.

Mapping the corpus. To facilitate access to the corpus metadata and visualization of the corpus works in their spatio-temporal context, each example was geocoded based on the location of the source work full unit's first public performance using ArcGIS and Carto. Whenever possible, the specific venue's geographical coordinates (latitude and longitude) were used; if venue could not be determined, the city's coordinates were used. By clicking on the premiere location, the user is able to access a portion of the metadata associated with each example as well as a score representation of the example and the randomly pre-selected recording of the source work full unit. Works' premiere locations can also be visualized in time based on a specific composer or genre or a sub-set of composers and genres. A sample map of the reserve data set with timeline animation is available here: https://epoudrier.carto.com/builder/adfbe0b9-093c-4eb7-b3ef-bb2b4b05a99f/embed.

Overall, it is worth noting that the most common premiere location for the full corpus is Paris, France (21.1% representation), with New York being the second most common (12.5%). On the timeline representation, there is a noticeable shift from Paris to New York when comparing the premiere locations from before to after the second world war, which may be related to several of the composers represented in the corpus having emigrated to the United States (e.g., Paul Hindemith and Arnold Schoenberg) at that time.

Corpus Characteristics

While the use of a corpus of examples collected by a third party shifts the sampling bias away from the researchers, generalization of findings will need to take into account the specific characteristics of the corpus. The following sub-sections present some descriptive statistics pertaining to the examples in the full corpus ($N = 719$).

Timeline. While most of the corpus examples were composed, premiered, and published in the time period from 1900 to 1950, there are a number of examples that fall outside of this period, with 22 works having been first completed before 1900 and 69 works having been first completed from 1950 onward. Thus, the actual period represented by the three phases in the creation of these musical works corresponds to a time range of 1877 to 2009 (see Figure 2).
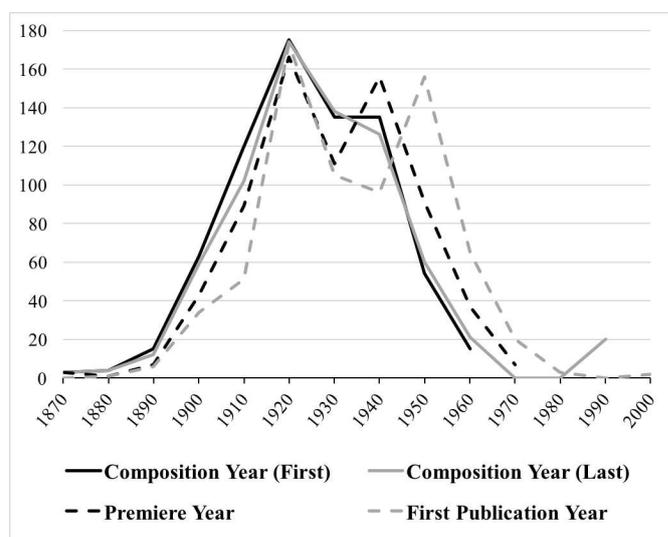

Figure 2. Distribution of corpus examples by decade based on composition year, premiere year, and first publication year. To account for composers' revisions of source works, metadata pertaining to composition year includes both first completed year and last completed year. Note that the spike in last completed composition year is represented by a single work, Olivier Messiaen's *Turangalîla Symphony*, from which 20 examples were extracted.

In general terms, the average first completed composition year represented by the corpus is 1929 ($SD = 16$) and the last completed composition year, which accounts for composers' revisions of the source works, is 1932 (SD = 18.6); average year for premiere and first publication are only slightly later, 1934 ($SD = 17$) and 1939 ($SD = 18.5$), respectively. However, while composition ($N = 719$), premiere ($N = 711$), and first publication ($N = 712$) followed each other closely, there appears to be a noticeable delay between premiere and first publication in the 1940s. Of special interest in this corpus is the potential disruptive effect of the two world wars (1914-1918) on the process of musical creation, from composition to premiere and publication, and this observation suggests that it may be a fruitful avenue of research.

Genre and sub-genre. The corpus source works were categorized into four genres based on the type of ensemble for which they were written: orchestral, instrumental, vocal, and

stage works. Overall, orchestral and instrumental works account for 38% and 36% of all source works, respectively, while both vocal and stage works account for only 13%. Each of the four genres was further characterized based on style and specific instrumentation (see Figure 3). When taking sub-genre into consideration, the most represented works are for chamber ensemble ($N = 175$), accounting for 24.3% of the corpus works (and extracted examples), and the least represented are ballets and theatrical works (e.g., Arthur Honegger's *Jeanne d'Arc au bûcher*, a dramatic oratorio for spoken roles, soloists, children's choir, mixed chorus, and orchestra, which was composed in 1935 and premiered in Basel, Switzerland in 1938).
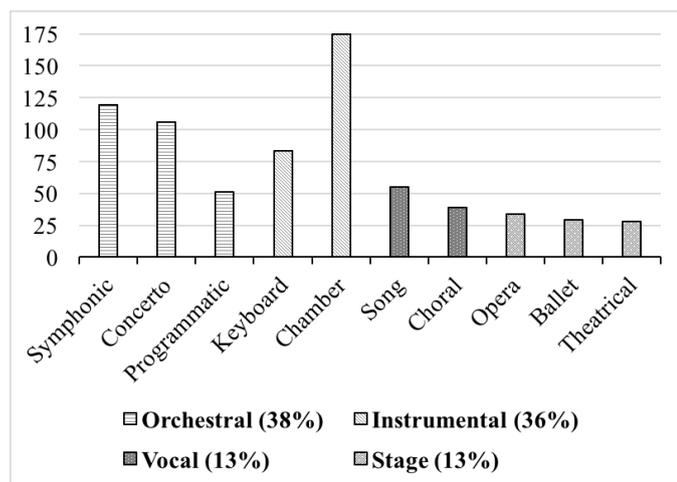

Figure 3. Genre and sub-genre representation. Percentage representation for each of the four genres is given, and sub-genre distribution is illustrated by the corresponding number of works in the full corpus.

Meter and tempo. Metadata pertaining to start time signature and tempo for both source work full units and extracted examples were collected. Overall, the most common start time signatures used are 4/4 and 3/4, which together account for 45.9% and 38.1% of source works full units and extracted examples, respectively. Given the wide range of time signatures used, these were further categorized based on the specific number of beats in the notated bar (duple, triple, quadruple, and odd, which feature uneven subdivision of the notated bar) or the relationship between notated beats and subdivision units (e.g., compound and polymetric); the type "none" refers to works in which rhythmic duration is notated without a reference to a time signature (see Figure 4).

In addition to quadruple and triple time signatures, duple time signatures (especially 2/2 and 2/4) are also relatively common, accounting for 19.3% and 22.2% of source work full units and extracted examples, respectively. In contrast, odd and polymetric time signatures are the least frequently used, respectively accounting for only 3.4% and 2.6% of start time signatures for source work full units ($N = 24$ and 18), and for 5.8% and 6.9% of start time signature for extracted examples ($N = 40$ and 48). However, while start time signature representation for source work full units and extracted examples are fairly similar overall, quadruple time signatures (which include 4/2, 4/4, 4/8, and 4/16 as well as 8/8 and 8/4) and polymetric time signatures (e.g., 4/2 and 3/4 superposed)

are somewhat more frequent at the beginning of the extracted examples than at the beginning of the source work full units. The relatively higher prevalence of polymetric time signatures at the beginning of the extracted examples is consistent with the nature of the extracted examples, all of which feature some form of polyrhythm, but the increased representation of quadruple time signatures has no theoretical basis and is worth exploring further. Finally, one aspect of the extracted examples that is not represented here is that 14% of these excerpts feature changing time signatures.
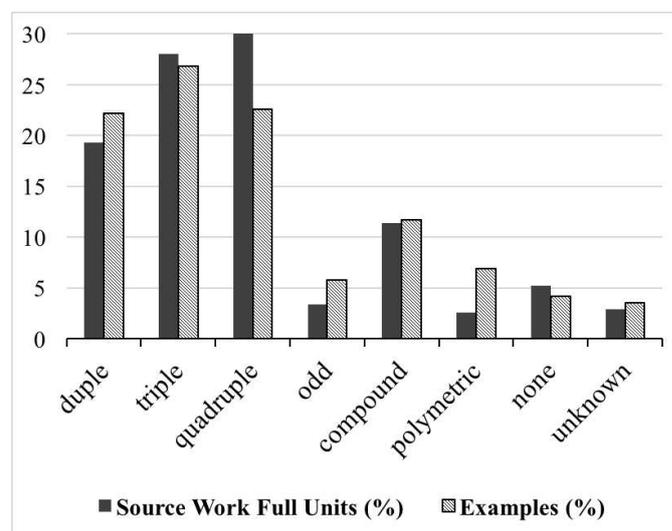

Figure 4. Percentage representation of start time signatures for source work full units and extracted examples based on metric type. Start time signatures could not be identified for 20 source work full units and 24 extracted examples, due to difficulties in locating a published score.

Start tempo metadata (expression, notated beat, and metronomic rate) were also collected. Expression markings (e.g., Allegro) are featured at the beginning of 93% of source work full units (based on metadata gathering at 97.2% completion), while specific metronomic rates for the notated beats are featured at the beginning of only 73% of the source work full units (and only 70% of the extracted examples, based on metadata gathering at 96% completion). When provided, the average notated start tempo is 95 bpm for the source work full units ($SD = 34.2$), while that of the extracted examples is 103 bpm ($SD = 37.0$), a difference that calls for further analysis. Because of the limited availability of metronomic rates as well as the difficulty involved in computing tempo based on the notated examples alone, the recorded duration of each example was also collected using randomly selected recordings from the Naxos Music Library.

Measuring Complexity

Using our reserve dataset, we performed a number of exploratory tests focusing on the interactions of tempo, genre, composer, year, and various aspects of rhythmic complexity, such as entropy and rhythmic variability. Specifically, we examined the composite rhythms of each excerpt, both as a single (global) feature and as a pair of interacting rhythmic groups (e.g. "3 against 2"). Three research assistants were tasked with placing parts ($N = 3$ to 30) into one of two groups;

examples that featured only two parts were automatically divided into two groups, and there was one example that featured a single part, which was excluded from the procedure. Research assistants were presented with the score and asked to assign each part to one of two groups based on how rhythmically and metrically similar the parts were to each other, with the parts that were least similar to each other having to be placed in different groups. These coders (all music graduate students) were largely in agreement. When two of the three graduate students were in agreement, the majority opinion decided the separation. In the rare instances where there was no agreement, we examined the score and made a judgment about the appropriate grouping. All of these cases were due to duplicate or missing parts. The analyses presented here therefore examine the single composite rhythm of excerpts as well as the interaction of two composite rhythms, each of which results from one of the two contrasting rhythmic groups within each excerpt.

Measurements

As a starting point, we decided to examine aspects of rhythmic entropy and variance. *Entropy* (a general measure of uncertainty and variability) was calculated as the Shannon entropy of the composite rhythm of the entire excerpt (the global composite rhythm), as calculated with the *Humdrum Toolkit* (Huron, 1995; Sapp, 2005; specifically, the *infot* and *beat* tools). This allows for a measurement of some degree of "unexpectedness" in the rhythm. Similarly, we calculated the *average variability* of each instrument, as calculated with the normalized pairwise variability index (see Daniele, 2017; Daniele and Patel, 2003; Grabe and Low, 2002). Rather than examining a composite rhythm, this measurement examines each rhythm individually, providing an average metric for variability.

Additionally, we examined the event density of each musical excerpt (as notated in the score), as well as the density in "clock time." *Event density* was calculated simply as the number of events per notated measure, whereas *clock time density* was measured on the performed time extracted from a randomly selected recording. Tempo was also calculated based on the recorded performance following these steps: the number of beats per measure were multiplied by the number of measures in the notated example, and the resultant number was divided by the time elapsed in the recording. For example, if 4 measures of 4/4 were performed in 5 seconds, the tempo was calculated as (4*4)60/5, or 16*12, resulting in an extrapolated tempo of 192 bpm.

Finally, in addition to examining these metrics in isolation, we also examined the difference between rhythmic groups, as represented by their composite rhythms.

Change over time. With the reserve dataset, it seems that there is no change in any of the aspects we examined over time, but a number of relationships point to possible avenues for future research. There was also no significant change in tempo, nPVI, entropy, or the difference in entropy and nPVI between rhythmic groups. Figure 5 indicates a gradual slowing of tempo over time in our reserve dataset, in keeping with much of the research on performance practice during this period (see Cook, 2013; Philip, 1992). These tempi, however, were deduced from contemporary performances of the pieces, meaning that rather than reflecting performance practice, they reflect a difference in contemporary approach from pieces composed in different periods.
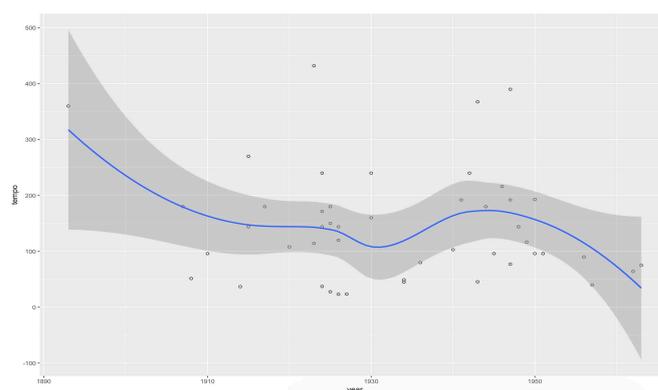
Figure 5. Performance tempi plotted against first composition year.

Interestingly, there seem to be a number of measurements in which there is a mid-century dip, including the entropy of the composite rhythm entropy (see Figure 6), and the difference in entropy between the two rhythmic groups. The time period over which decreasing entropy can be observed seems to correlate to the period during which a delay between premiere and first publication year was previously observed. As noted earlier, as this period also corresponds with the emergence of the second world war, this is a result that calls for closer examination.
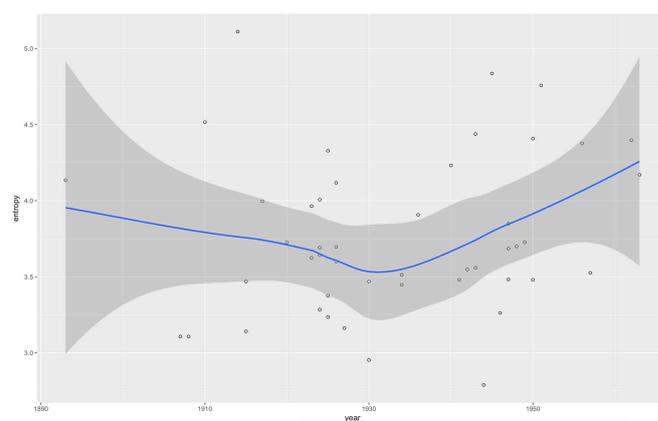
Figure 6. Entropy of the composite rhythm over time.

Predicting genre. The interaction between rhythmic characteristics and genre is also quite interesting. Figure 7 presents the variable importance of certain characteristics when used to train a generalized linear model classifier (measured here with the absolute value of the *t*-statistic of each feature) to predict genre. Although the classifier itself performed significantly better than chance (71% accuracy), it would likely do much better given more data. These results suggest that the difference in entropy between rhythmic groups plays a large role when predicting vocal works, whereas the difference in nPVI between groups is possibly more predictive of instrumental works. Further analysis is needed to pick these points apart, as the two measurements likely contain a certain amount of covariance. Finally, tempo and event count also appear to be important in predicting orchestral works, which similarly warrants further inspection.
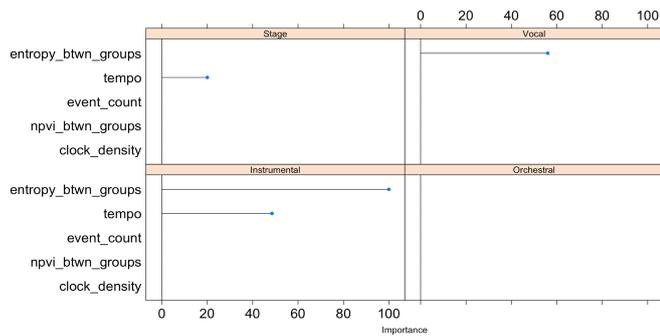
Figure 7. Log Odds Ratios resulting from a mixed effects logistic regression model predicting metric type. None of the features were significant predictor.

Conclusion

So far, claims about a significant increase in rhythmic complexity over the 50 years represented by the Suter (1980) corpus have not been substantiated, at least from the perspective of global measures of variability (entropy and nPVI) and density applied to a stratified sample (reserve dataset) of examples that feature a wide range of polyrhythmic structures. This may be understood to attest to a certain stylistic homogeneity in the music of this period, although this observation remains to be tested. Nonetheless, some trends based on national origin and genre have been identified, and the wide variance observed in measures of rhythmic regularity based on composer, nationality, genre, and pre- vs. post-war composition year suggests that there are other factors at play, which future analytical work will address.

Although there were no significant differences observed, it would seem that there are a number of avenues in which further research would be warranted. Firstly, polyrhythmic gestures change over time, but not necessarily in a linear way. Our preliminary results suggest that it is likely that the middle of the twentieth century differs from earlier and later periods in terms of polyrhythmic practice. Secondly, it would seem that the types of features used in these polyrhythmic compositions might be genre-dependent. Future work would ideally explore to what extent this might be the case, what other features are involved, and how the changing nature of complex rhythms might inform our understanding of twentieth-century music.

References

Cook, N. (2013). *Beyond the score: Music as performance*. New York: Oxford University Press.

Daniele, J. R. (2017). A tool for the quantitative anthropology of music: Use of the nPVI equation to analyze rhythmic variability within long-term historical patterns in music. *Empirical Musicology Review*: EMR. Retrieved from http://emusicology.org/article/view/4893

Grabe, E., & Low, E. L. (2002). Durational variability in speech and the rhythm class hypothesis. *Papers in Laboratory Phonology, 7*, 515-546.

Huron, D. (1995). *The Humdrum Toolkit: Reference manual*. Menlo Park, CA: Center for Computer Assisted Research in the Humanities.

Huron, D. & Ommen, A. (2006). An empirical study of syncopation in American popular music, 1890-1939." *Music Theory Spectrum, 28*(2), 211-232.

Krebs, H. (1999). *Fantasy pieces: Metrical dissonance in the music of Robert Schumann*. New York, NY: Oxford University Press.

Large, E. W., & Jones, M. R. (1999). The dynamics of attending: How people track time-varying events. *Psychological Review 106*(1), 119-159.

Patel, A. D., & Daniele, J. R. (2003). An empirical comparison of rhythm in language and music. *Cognition, 87*. https://doi.org/10.1016/S0

Philip, R. (2004). *Early recordings and musical style: Changing tastes in instrumental performance, 1900-1950*. Cambridge University Press.

Poudrier, È. (2009). Local polymetric structures in Elliott Carter's *90+* for piano (1994). In B. Heile (ed.), *Modernist legacy: Essays on new music* (pp. 205-233). Farnham, England: Ashgate.

Roeder, J. (1994). Interacting pulse streams in Schoenberg's atonal polyphony. *Music Theory Spectrum, 16*(2), 231-249.

Sapp, C. S. (2005). Online Database of Scores in the Humdrum File Format. In *Proceedings of the International Society for Music Information Retrieval* (pp. 664–665).

Suter, L-M. (1980). Les polyrythmes dans la musique de vingt compositeurs de la première moitié du vingtième siècle, 1900-1950 [Polyrhythm in the music of twenty composers from the first half of the twentieth century, 1900-1950]. (Unpublished doctoral dissertation). University of Bern, Bern, Switzerland.

Temperley, D. (2004). *The cognition of basic musical structures*. Cambridge, MA: MIT Press.

Volk, A. (2008). The study of syncopation using Inner Metric Analysis: Linking theoretical and experimental analysis of metre in music. *Journal of New Music Research, 37*(4), 259-273.