Theoretical Corpora and Compositional Practice

JOSEPH PLAZAK [1]
Illinois Wesleyan University

ABSTRACT: By combining two distinct research tools, algorithmically generated musical corpora and computational models of music perception, Dean & Pearce (this volume) report a novel methodology and results that might be of interest to both composers and theorists of non-tonal music. This commentary briefly elaborates on one particular implication of their work, namely, the potential to predict the salience of theoretical principles prior to implementing them within compositional practice. Several limitations specific to this particular implication are discussed.

THERE are many examples in which theories of music follow the implementation of compositional processes, such as applying Roman numeral-style harmonic analyses to the music of J.S. Bach, or utilizing set theory to understand the works of the Second Viennese School. Despite a few instances to the contrary, this arrangement, in which compositional practice tends to precede standardized theories, can be observed in many types of music. One important implication from Dean & Pearce's (2016, this volume) work using algorithmically generated corpora is the possibility of testing abstract compositional principles (i.e., compositional theories) via cognitive models of auditory expectation prior to compositional implementation. Specifically, through representative musical corpora and perceptual models, it becomes possible to predict the salience of certain musical features, and these theoretical predictions could inform, rather than follow, compositional practice. Therefore, knowledge of the methodology employed by Dean & Pearce should be of great interest to both composers and improvisers engaging in new musical practices. However, the utility of this methodology is not without limitation, as it is highly dependent on both the quality of the corpora that are employed and the accuracy of the chosen perceptual models. While the authors are careful to discuss some limitations imposed by both their algorithmically derived corpora and their auditory model, a few elaborations are offered here.

CORPORA & MODEL LIMITS

While it is true that compositions using serialized pitch patterns could resemble the pieces within Dean & Pearce's algorithmically generated corpora (both tonal and non-tonal), there are many musical nuances that are lost. This should not be taken as a criticism against methodical simplification, but rather as a caution towards what can actually be gleaned from analyzing "sequential pitch structures" as opposed to melodies. While serial music is often associated with the compositional goal of avoiding pitch centricity (i.e. the perception that one or more tones are more prominent than the others), there are still plenty of ways in which composers artfully explore pitch centricity within serialized compositional methods. Common musical structures that create a sense of pitch centricity include musical stresses such as: agogic stress, register stress, metric stress, dynamic stress, articulation stress, timbre stress, or repetition stress. Even though a composer working with serialized structures is likely to begin by choosing a well-formed pitch row and then exploring the standard permutations of this row, these are only the first of a seemingly limitless number of compositional choices. In seeking to understand what an algorithmically generated "sequential pitch structure" corpus could tell us about the perception of actual non-tonal music, it might be helpful to first examine a serialistic musical work with all musical features, including features of rhythm, harmony, timbre, dynamics, etc., and then attempt to determine the extent to which pitch patterns explain the total variance within a given perceptual model. For example, one might elect to analyze a given piece with strict serial pitch patterns (e.g., Krenek's op.84), train a perceptual model (e.g. IDyOM) using the actual row matrix chosen by the composer, and thereafter determine the extent to which this training

reduces information content when compared to a model with more nuanced musical parameters. Such an approach might serve to calibrate and contextualize the utility of algorithmically generated corpora of "sequential pitch structures," like those employed by Dean & Pearce (this volume).

Beyond discussing their corpora, the authors also briefly discuss the limitations of the IDyOM model. Despite IDyOM being a monophonic model (Pearce & Wiggins, 2012), Dean and Pearce attempt to model polyphony by utilizing what might be considered an "arpeggiation" technique. Namely, simultaneously sounding pitches are entered into IDyOM sequentially from the lowest sounding note to the highest sounding note. The authors do not draw strong conclusions based on their polyphonic models, so there is no need to dwell on the limitation of this strategy. However, it does bear reminding that even within purely serialized music, subsequent pitch-related expectations can be made relative to a perceived pitch center, rather than being based on absolute pitch/pitch-class intervals. Such would be the case when a given 12-tone pitch row is performed above a sustained drone tone; not all pitches should be assumed to be perceptually equivalent. I mention this not so much as a criticism of the authors' methodology, but rather as a reminder that there is still much work to be done towards refining perceptual models of music, and in particular, polyphonic music.

ANALYSIS

After describing their corpora and perceptual model, results in the form of "mean information content," are provided for each corpus under a variety of training conditions. Throughout the results section, the data are distilled and reported as this single value. In the absence of other relevant information, such as, standard deviations, degrees of freedom, and/or confidence intervals, it is difficult to objectively assess the significance of such small mean IC differences (e.g., 0.01 differences between mean IC scores). According to the authors, these reported results reflect an average of 10 mean IC values (Dean & Pearce, 2016, p.8), and while perhaps verbose, including the raw data for these averaged-averages might prove to be beneficial for those unfamiliar with IDyOM modeling, or those who might consider using these results in future meta-studies. Further, there are several places where the authors allude to the statistical effect size of their results: "these values are much higher" (p.10); "dramatically lower IC" (p.10); "substantially predictive" (p.17); and "corpus benefitted slightly" (p.19), to name a few examples. Such claims are difficult to objectively assess from raw mean scores alone, especially for musicians who might consider making compositional decisions based on these results.

One of the main claims of this paper was that "algorithmic serial melodic (monadic) compositions are useful in information-theoretic modeling of pitch sequences in a range of non-tonal music" (p.17). In the context of re-examining the relationship between compositional practice and theoretical inquiry, this is an exciting finding. The authors also state that, "even the algorithmic serial tonal melodic corpus provided some predictive power, though much less so than non-tonal melodies, as hypothesized" (p.17). Based on the reported data in Tables 1 & 2 (p.10 and 11, respectively), this pattern appears to only hold true in one case, which involved modeling the Tonal Polyphonic Dataset after pre-training with non-tonal pieces. In that case, training with Tonal Polyphonic pieces 7-12 resulted in a mean IC of 5.35, whereas training the same dataset with the Non-Tonal Polyphonic pieces had more predictive power (mean IC = 5.29). However, during all other reported iterations of training and testing, tonal corpora seemed to produce the same, or lower, mean IC scores as non-tonal corpora training. The issue at hand seems to largely be a matter of definition. In some cases, the authors report the results of tonal, non-tonal, and both tonal & non-tonal training combined (i.e., Figures 1 & 2), whereas in other cases, only non-tonal and combined tonal & non-tonal trainings are reported (i.e., Tables 3 & 4). With regards to the latter instances, an increase in mean IC when "adding" the tonal training set to the non-tonal training set seems to be used as evidence that the non-tonal training was more effective, however, without actually reporting data from the Tonal Corpus training condition, it is difficult to assess the claim that tonal training is relatively less useful for a given dataset. This point is raised because it would be particularly interesting to see the influence of algorithmically generated tonal sequence training on the mean IC results for non-tonal datasets, and the authors themselves seem to have predicted this potential influence (p. 3). These results (from models trained with algorithmically generated tonal pitch structures) would provide valuable insight not only into how IDyOM modeled the reported Webern, Schoenberg, and piano improvisations, but again, also to those who might be interested in using these results to make informed compositional decisions.

SUMMARY

As implied by the title of their article, "Algorithmically-generated corpora that use serial compositional principles can contribute to the *modeling of sequential pitch structure* in non-tonal music" (italics added for emphasis), the authors report on a specific element of non-tonal music perception, and are careful to avoid holistic claims. Given the so-called "complexity" of this type of music, and the complexity of music perception in general, sequential pitch structure is a reasonable musical aspect to launch this type of inquiry. The methodology presented in this paper could easily be extended to other musical features (i.e., rhythm, texture, orchestration, etc.) and other musical styles (i.e., electronic music, popular music, jazz, etc.); it is my hope that scholars continue to explore these possibilities. In order to bear utility for composers and improvisers, there remains much work to be done in order to overcome limitations currently imposed by existing musical corpora & music perception models. These caveats notwithstanding, and as evidenced by their paper, Dean & Pearce have already demonstrated new possibilities for re-examining the relationship between musical theory and compositional practice.

NOTES

[1] Correspondence can be addressed to: Joe Plazak, Illinois Wesleyan University, jplazak@iwu.edu.

REFERENCES

Dean, R. & Pearce, M.T. (2016). Algorithmically-generated corpora that use serial compositional principles can contribute to the modeling of sequential pitch structure in non-tonal music. *Empirical Musicology Review, 11*(1), 27–46. http://dx.doi.org/10.18061/emr.v11i1.4900

Pearce, M. T., & Wiggins, G. A. (2012). Auditory expectation: The information dynamics of music perception and cognition. *Topics in cognitive science*, *4*(4), 625-652. http://dx.doi.org/10.1111/j.1756-8765.2012.01214.x