SCHEMATIZING THE TREATMENT OF DISSONANCE IN 16TH-CENTURY COUNTERPOINT

Andie Sigler
School of Computer Science,
McGill University;
Computing Music

Jon Wild
Schulich School of Music,
McGill University
wild@music.mcgill.ca

Eliot Handelman
Computing Music

ABSTRACT

We describe a computational project concerning labeling of *dissonance treatments* – schematic descriptions of the uses of dissonances. We use automatic score annotation and database methods to develop schemata for a large corpus of 16th-century polyphonic music. We then apply structural techniques to investigate coincidence of schemata, and to extrapolate from found structures to unused possibilities.

1. INTRODUCTION

We develop a set of schematic dissonance treatments (i.e. schemata under which the uses of dissonance are classifiable) using a large corpus of mass movements (almost 1000) of Palestrina and Victoria, dating from the 16th century. Palestrina in particular has a resonance through the history of music as one whose style was raised to the status of a didactic norm. [1] As a result, Palestrina's practice (or a simplification of it) has been well known and imitated for centuries among academics and music students. [2] As a foil for Palestrina, we compare masses by Victoria, roughly contemporaneous and with a similar dissonance treatment. The wealth of available literature on the dissonance practice of this style gives us a departure point for developing a computational platform for its investigation, with a view to generalization.

[1] As pointed out in Alfred Mann's 1991 forward to Jeppesen's *Counterpoint* [2], one of several classic texts on the Palestrina style – as the title shows, the name Palestrina is all but synonymous with certain aspects of basic musical organization – in particular the way a "point" (i.e. a note – or perhaps a musical "idea") sounds and moves "counter" to (i.e. in relation to) another point or set of points.

[2] Including e.g. Haydn, Mozart, Beethoven, Schubert, Rossini, Chopin, Berlioz, Liszt, Brahms, Bruckner, R. Strauss, and Hindemith, who all are known to have used Fux's *Gradus ad Parnassum*, based on the Palestrina style ([1], Mann's introduction).

2. METHODOLOGY

2.1 Automatic Score Annotation

Part of our methodology for investigating dissonances is to look at *automatically annotated scores*. Seeing annotated scores helps us evaluate the correspondence of our specification to our intention, and develop to new schemata. It allows us to identify musical factors that would likely not have been apparent otherwise (i.e. in a situation where data was displayed in a musically non-intuitive way, or where scores had to be painstakingly scrutinized to locate scarce occurrences).

Using a web-based music analysis system produced by Computing Music, we generate annotated scores *on-demand* (at load time). It's possible to load and analyse any score (including ones outside the corpus under investigation), to load a random score from the corpus, or to "spin" through a corpus with a search for instances of a particular configuration, such that one keystroke displays a new annotated score focussed on the relevant measure, bringing together similar occurrences from disparate locations.

2.2 Saving Features

On a first pass through the score, we save a set of *features* for each dissonance, including duration, surrounding melodic intervals, metric weight, and type of attack, as needed to define our schemata (– as we developed and added new schemata, an initial set of features was expanded). Any features not used in a given schema are open to any value.

Saving a set of features for each dissonance rather than just applying a set of schematic filters on the first pass through the score has certain advantages. Suppose we run all dissonances through a set of filters, and several of them are labelled P for passing. Now if we want to ask questions about the set of passing notes (in fact matched by several different but related schemas) – e.g. how many are going up or down, or how many are half-notes – we have to *re-ask* some of the same questions we already asked in order to label them in the first place. As well, if we have a set of remaining *unlabelled* dissonances, we will have no idea *how* they failed the tests for the different labels, or what subsets of unlabelled dissonances might have in common.

We save feature-sets (and schema labels) in a database, so that we can query them in different ways; database searching also helps us develop schemata based on featural similarities of unschematized dissonances.

3. DEFINITIONS AND SCHEMATA

3.1 Dissonance and Meter

As in standard practice, we define the *dissonant intervals* as the minor and major second, perfect fourth, tritone, and minor and major seventh, and their compounds (i.e. with additional octaves). A *dissonance* occurs when two notes coincide or overlap in time and form a dissonant interval.

From the initial set of dissonances in a score, we remove certain fourths and tritones that participate in sonorities considered consonant. [3] If a perfect fourth is accompanied by an additional voice sounding a third or fifth (or their octave compound) below its lower note, it is considered consonant. Likewise a diminished fifth accompanied by the pitch a major sixth below its lower note, or an augmented fourth accompanied by the pitch a minor third below its lower note.

Meter can be thought of as a temporal grid. We generalize to metric *weight*, where different places in the measure are said to be equally "strong" or "weak." The downbeat is the strongest, followed by the divisions into halves, then divisions into quarters, then eighths.

The meters under consideration are *duple*, using whole-note divisions (e.g. 4/2; 2/2), or *triple*, using dotted-whole divisions (e.g. 3/2, 6/2, 9/2). We don't differentiate between whole notes in a duple meter or or dotted wholes in a triple meter; each one represents an equal beat.

3.2 One- and Two-voice Schemata

Although a dissonance is defined as a relationship between two notes in two different voices, commonly only one note needs to be "explained." [4] Typically, when one note is struck and then sustained (or reattacked) while the second note is struck (an *oblique* motion), and the second note is on a "weak" metric position, only the *second* note needs explanation, since the dissonance only occurs once the second note enters – we call the second note the dissonance (with respect to the first note). In these kinds of cases, we can schematize a dissonance treatment with respect to features of the voice containing the dissonance, and not the voice against which it dissonates. For example, a *passing* note is schematized by either of two different melodic shapes: (step up, step up) or (step down, step down). [5] It is simultaneously schematized by one of four different metric shapes: a half note on a weak half preceded by a duration of at least a half, a quarter note on a weak quarter, an eighth on a weak quarter, or an eighth on a weak eighth. We have defined several other "single-voice" schemata; these are summarized in Table 1.

A *suspension* is a *two-voice* schema, involving the suspended note as well as its counterpart, the "agent." The agent, or active voice, is an obliquely struck dissonance on a strong beat, after which the *other* voice (the suspension) is constrained to resolve downward by step. Since we've already set up machinery to find *attacked* dissonances, rather than to find notes that are dissonated against at a particular place in their duration, it's convenient to start the schematization of suspensions with the agent, rather than with the suspension note itself. When we find an oblique dissonance on a strong beat, we can pull in a feature set for the note against which it dissonates, and check whether the combination constitutes a suspension. Further description of suspensions can be found in Table 2.

3.3 Extending the pairwise model

We originally defined dissonances as occurring between *two* voices. One exception to this model that we have already addressed is the consideration as consonant of fourths and tritones that are covered by certain notes in a lower-sounding voice – these are "vertical" or *harmonic* schemata. Apart from these, we have so far used a pairwise model to schematize dissonances between any two voices. But we found we had to extend the pairwise model to account for some dissonances. These are summarized in Table 3.

1. We find that if a note is consonant with an agent of a suspension, it can be dissonant with the suspension without further constraint; as well, we find situations where a note is dissonant with an agent, but explicable as consonant with the suspension.

2. On a weak quarter, two quarter notes or eighths (or one of each) may be dissonant with respect to each other, if there is a third voice such that each is explained as consonant, passing, neighboring, a cambiata, or an anticipation with respect to the (same) third voice. (See Figure 1.)

3. A note m that is dissonant within a given pair of voices is in condition M if it has the same pitch class as a note in a third voice that was already sounding when m entered, and is sustained at least until the end of m. Notes in condition M are often approached and left by leap. A note in condition M may be attacked simultaneously with a dissonance; in this case the note not labelled M will be explained (e.g. as a passing or neighbor note) with respect to the third voice. [6]

4. DISCUSSION: EXCEPTIONS AND INDUCTION

At the time of this writing, there are still ∼360 dissonances in the Palestrina-Victoria corpus that are not explained by

[3] These correspond to major and minor triads in root position or first inversion, and diminished triads in first inversion, though these designations are anachronistic for the 16th century.

[4] Informally, *explaining* means locating a theorized schematic dissonance treatment to which a dissonance corresponds.

[5] I.e. (step up, step up) gives a figure of *three notes*, including a step up *to* the dissonating note called the "passing note", and another step up *from* the passing note.

[6] In fact, if we look at half notes that are dissonant counterparts to condition M, we find that they are *all* passing notes, with six exceptions that are upper neighbors – and these six are all in the same mass of Victoria. This is an example of a unique dissonance treatment, used motivically, that is clearly related to the more common passing version.

Symbol	Name	Melodic schema	Metric schema	Attack
P	Passing	(step up, step up) (step dn, step dn)	weak quarter eighth on weak quarter weak eighth weak half after $\geq$ half	oblique
N	Neighbor	(step up, step dn) (step dn, step up)	(same as for P)	oblique
C5,C4,C3	(5/4/3-note) Cambiata	(step dn, third dn, step up, step up) – or first *n* notes of this	weak quarter	oblique
A	Anticipation	(step, repeat)	weak quarter weak eighth	oblique
E	Echappée	(step up, leap dn) (step dn, leap up)	weak quarter	oblique
F	"Fake" suspension	(step or repeat, step dn)	syncopated whole syncopated half syncopated dotted-half	oblique
Q,Qx	Third quarter	(step dn, step dn)	quarter on weak half; Q if after $\geq$ half, otherwise Qx	oblique simultaneous
L	Leap of third	(third dn, step up)	weak quarter	oblique

Table 1: "Single-voice" Schemata

Symbol	Name	Description
S,G	Suspension, Agent	Suspension S is sustained or reattacked on the same note; agent G strikes oblique dissonance; S moves down (to its *resolution*) by step on a weaker beat than G.
T,T2,G	Suspension with third-skip, Agent	As S, but with resolution (third dn, step up) in quarter notes; the note skipped down to can be dissonant (called T2).

Table 2: Suspensions

Symbol	Name	Description
Gc	Consonant with Agent	Dissonant with a suspension S or T but consonant with its agent. Or dissonant with a "Fake" suspension F and consonant with its "agent."
Sc	Cons. with Suspension	Dissonant with an agent but consonant with its suspension.
M/M2, Mx	Match	Has the same pitch / pitch class as a note in a *third* voice already sounding when M entered, and is sustained at least until the end of M. M's dissonant counterpart Mx is attacked simultaneously with M; explained as P or N with the *third* voice.
W	Weak-quarter clash	On a weak quarter, a dissonance between two quarter notes or eighths (or one of each), such that each is consonant, passing, neighboring, a cambiata, or an anticipation with respect to some (same) third voice.

Table 3: Schemata: Extending the Pairwise Model

any of our schemata (versus $\sim$194100 that *are* – we've successfully schematized $> 99.8\%$ of dissonances in the corpus). Unschematized dissonances are marked with an X in annotated scores. There are quite a few errors (e.g. wrong notes or durations) in our corpus, and it looks like a considerable proportion of Xs are due to these. The ability to quickly navigate to problematic dissonances allows us to make corrections where they are necessary (i.e. by comparison with another edition) – correction of the corpus is currently underway. This method doesn't locate *all* errors in the corpus, but it does point out especially "bad" ones from the point of view of dissonance.

Examining Xs is also part of our development methodology for formulating new dissonance categories. For instance, by doing some filtering on a database of unmatched feature-sets, we noticed that there were 54 unschematized dissonances that are on weak quarters and are approached by a third down and left by a step up. We wrote this schema into our specification ("L" in Table 1), and then were able to "spin" through the instances in the corpus to see whether the schema met our expectations on the annotated scores. [7]

In another database exploration case, we began by observing that there were quite a few unschematized dissonant half- and quarter-notes on beat one, which were approached and left by a step. This preliminary schemati-

[7] After finding this dissonance in the database, we observed that it is mentioned (as possibly an "archaism") in [3], p. 220. Jeppesen's study proves to be a tour de force of detail – for example, on p. 268 he shows a suspension that jumps down a fifth before leaping back up a fourth to its resolution, saying that as far as he's observed, "this occurs but once in the whole collection of Palestrina's compositions," despite being a standard practice in [1]. We don't find a second occurrence in Palestrina, and it occurs once in our Victoria corpus.

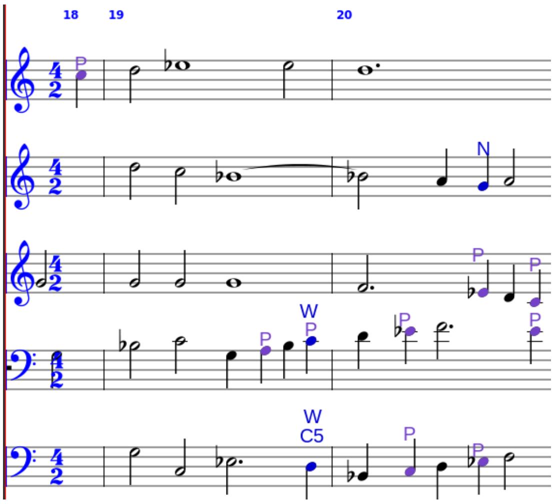

Figure 1: W (weak-quarter clash) are "explained" dissonances (or consonances) with respect to other voices (in this case P and C5), but simultaneous, unclassified dissonances with respect to each other.

Figure 2: The tied note in the third voice marked "F S" is a fake suspension (F) with respect to the bottom voice, and a "real" suspension (S) with respect to the agent (marked G) in the voice above.

zation was obviously too general to keep as a final labeling, since we don't wish to allow passing and neighboring notes on strong beats indiscriminately. But looking through these instances showed us that (along with a small number of less explicable occurrences), there were a couple of schematic situations. One such situation occurred when the dissonance in question had an *agent* as its dissonant counterpart, while being *consonant* with the corresponding *suspension* (shown as label Sc in Table 3). We also were able to refine our definition of *suspensions* by looking at these unschematized strong-beat dissonances. Our original definition stipulated that the agent must be *consonant* with suspension's note of *resolution* (whether or not the agent is still sounding at the time of the resolution). In fact, we find there is one situation where this rule doesn't hold: when the suspended interval is a diminished fifth, resolution forms a fourth (– dissonant) with the

Figure 3: A unique structure of simultaneous dissonances: two passing notes, a neighbor, and a cambiata. (Palestrina: Laudata Dominum, Gloria)

agent. When this happens, the agent always moves up a step to meet the resolution in a (consonant) *third*.

When exploring for new schemata, we sometimes come against occurrences that are interestingly rare. For instance, we find that there are six third-quarter passing notes going *upward* in the combined Palestrina-Victoria corpus. Of these, four are in one mass of Victoria, and are essentially repetitions of the same single situation. The remaining two are separate instances in Palestrina. These kind of instances open musicological questions as to the interpretation of these scare occurrences: *why* was this possibility used just here, and practically nowhere else.

Database exploration works not only for induction of new schemas, but for deeper exploration of defined schemas. For instance, if we look at the feature set for the relatively rare half-note lower-neighbors, we find that *most* of them (in Palestrina 119/150, or 79%) are a perfect fourth above the note they dissonate with. A few (13, or 9%) are a tritone below, and on closer inspection, these all seem to take part in very similar cadential figures. Victoria uses the tritone/cadential lower neighbor somewhat more often – 20/83 or 24%, and the perfect fourth above 43/83 or 52%: a similar but less dramatic tendency.

Likewise, we find that our category for "fake" suspension (F) (which Jeppesen calls a "consonant fourth") never occurs with a tritone, and in fact *always* occurs with either a fourth, or (less frequently) a fourth *and* seventh or second (i.e. with respect to two different voices) . Furthermore, the F which is *only* a fourth at its onset is almost always accompanied by a suspension (S) of a seventh or second on the next strong beat (Figure 2) – the fake suspension of a fourth with *no* seventh or second at all is found only 16 times in the Palestrina corpus and never in Victoria. We could continue this line of musicological investigation by surveying for further details, finding e.g. those fake suspensions which are a half note in duration, or those introduced by leap, or those which include a dissonance of a *minor* second or *major* seventh (a rare occurrence), or those which have a resolution of a *major* second (relatively rare).

We wonder: would it be feasible to *automatically* induce dissonance treatments over a corpus (i.e. start from scratch and have a search deliver a set of schemata that are used a minimum number of times in a corpus). Although this would be computationally expensive, it seems possible.

The strategy for doing so, however, is not completely transparent. If we address the subset of one-voice schemata, we can imagine trying to cover the set of dissonances with minimal explanatory schemata (with the heuristic that more proximate intervals have to be part of a schema before more distant intervals can be included). For conjunctions, this is straightforward enough (e.g. must be on a weak quarter *and* resolve down). For disjunctions, we would have to infer whether a reduced set of features should be specified, or whether to use a wild card. We would have to be careful not to overfit schemata, which would result in a large number of highly specific schemata instead of a smaller number of more general ones (e.g. a passing note figure, once completed can be followed by a step up, or a leap up, or a third down, etc.). There's also no obvious way of joining multiple discovered schemata under one descriptive tag. For example, eight different schemata emerge for what we call "passing notes" (depending on their position, duration, and orientation) – and this is not including third-quarter passing notes, which we have chosen to name differently.

The schemata found would be constrained to be described by the feature set we're examining. We've used shorthand features such as "weak quarter," generalizing second and fourth quarters, and "leap up," generalizing several intervals. If we started off an automatic schema induction with these generalized features, it would be powerless to differentiate them (– generalizing *reduces* our power as human experts to differentiate them, but we still stand a chance of doing so by looking at scores). On the other hand, if we start with a *larger* feature set, we increase the search space exponentially, but add an interesting layer of *feature* induction. Even if we start with a larger feature set, we're still constrained by pre-process feature selection, whereas humans are free to add features midstream.

We won't discuss here the added problem of trying to induce two-voice schemata such as suspensions from scratch, nor the various three-voice schemata. We would also need to consider *harmonic* treatment: dissonances may be treated differently when they're a part of a chord (aside from the chords we have already discussed, for some corpora seventh chords, root-position diminished triads, or second-inversion triads have special status). Having errors in the corpus also complicates the picture.

While automatic schema induction is an interesting concept, for the time being it seems that using database queries and automatic score annotation to facilitate deep interaction of human intelligence with a musical corpus is still the most effective procedure.

5. STRUCTURING DISSONANCES

So far, what we've described are specific filters defined on feature vectors. These filters assign tags to notes, labeling the dissonance treatment of the note. Now we have the opportunity to see how these dissonance treatments interact. For instance, it's quite common to have two or more passing notes in different voices at the same time. What other combinations of dissonance might occur? For this analysis, we don't have to develop new schemata and filter for them, we merely have to *build structures* out of the dissonances we already have.

The procedure is this: we take a set of labeled dissonances, and build graphs of temporal relations between them. For the purpose of this example, we keep the space small by only examining a subspace of temporal relations between dissonances. We use three types of temporal relation: monophony (i.e. one or more notes beginning and ending at the same time), inclusion (i.e. a note's duration being *within* the duration of another note), and overlap.

We also use a subset of dissonances: passing and neighboring tones, third-quarters, anticipations, échappées, cambiatas, dissonant leaps of a third, "real" and "fake" suspensions, and weak-quarter clashes. The experiment reduces each score to *just* the notes marked with these labels, and then constructs polyphonic structures out of the remaining subscore. That is, we will connect tagged dissonances that are in temporal contact with one another, then examine sets of connected dissonances in the corpus. In what follows, we are counting not *notes*, but *structures*, which can contain one or more notes.

We obtain 297 different structures by this method – 243 in Palestrina and 175 in Victoria, with 121 in their intersection, and therefore 122 in Palestrina but not Victoria and 54 in Victoria but not Palestrina. [8] Of the 297, 113 occur only *once* in the combined corpus, while another 80 occur fewer than five times. In general, we see a relatively small number of structures occurring very frequently, and a large number of structures occurring rarely.

The most common structures in Palestrina and Victoria differ only slightly. The most frequent for both composers is the lone passing note, followed by the suspension, the double-passing note (i.e. two simultaneous passing notes), and then the (lone) neighbor. The next most common for Palestrina is the third-quarter passing note, then simultaneous passing and neighbor notes, and simultaneous neighbor notes. Victoria would be the same, except the third

[8] The absolute numbers themselves are not of great interest, and we don't offer a proper statistical analysis, we only mean to give a general orientation as to the structural variety available from the point of view of this experiment. The numbers are, furthermore, provisional since we're still correcting the corpus, but the great majority of rare structures are *not* due to corpus errors.

quarter dissonance is slightly more rare in Victoria, appearing after the latter two.

Other structures show a greater difference in practice between the composers. For instance, just looking at structures with double suspensions, we find that Palestrina resolves these at different times (i.e. one resolution coming a quarter note before the other) over half of the time, whereas Victoria only resolves them at different times about 10 percent of the time. [9] We find also that simultaneous "fake" suspensions don't occur in Palestrina, while there are 26 instances in the Victoria corpus. A figure in which a note is a dissonant third-quarter with respect to one voice at the same time as being interpreted as the agent of a diminished suspension [10] in another voice is found 23 times in Palestrina and once in Victoria. The cambiata occurring within the duration of a suspension, and the double third-quarter dissonance are also much more frequent in Palestrina than in Victoria. Everything found in Victoria more than three times is also found in Palestrina at least once – it's not obvious if this is an artifact of the difference in the sizes of the corpora, [11] or whether it reflects on the practice of the composers.

6. EXTRAPOLATION AND NEW STRUCTURES

The distribution of structures of labelled dissonances, with many structures used only once or a handful of times, shows us that we are not dealing with a closed set of reusable possibilities, but a *composable* space. This suggests that it's possible to build structures that are not in the corpus, but that are within the matrix of possibilities outlined by the corpus. In efforts to build style-copying automata, a trend has been to *re-use* and *re-combine* elements found in a corpus. But since it is the responsibility of the artist to offer *something new* in each work, reasoning about *unused* structures is essential for deeper exploration of corpus extension.

What we present in this section was not constructed automatically; we simply show that the structures we obtained from labelled dissonances seem to constitute a set with missing elements which *might* have been used in the corpus. It is our opinion that it would be possible to construct these automatically, and that in any case, the set of unused possibilities (and the set of once-used possibilities) are an avenue of insight into the nature of composition. Our ability to schematize the treatment of a great majority of dissonances in the corpus points to a constrained and rule-bound composition practice. How does this relate to the obligation to create new and different works? And is it possible to *reason* about newness and difference? In this section we suggest an approach.

We can proceed rather conservatively: instead of trying to invent complex and exotic new combinations that might be realizable, we can start by looking for unfilled niches that are relatively simple. For instance, if we take the subset of structures that consist of more than two *simultaneous* dissonances including at least one cambiata and one neighbor, we find *a single structure* that occurs once: a cambiata, neighbor, and two passing tones at the same time. This means that the *simpler* cambiata, neighbor, and *one* passing tone never occurs! We also see other obvious combinations including a cambiata and two neighbors, two cambiatas and a neighbor, and a cambiata, two neighbors and one passing note. It is simple to enumerate all of the possibilities in this small combinatorial space. [12]

For a given constructed dissonance structure, it's not guaranteed that it is realizable. We can try to realize it systematically by generating and testing candidates. The space of candidates is small enough to be tractably enumerable, especially if we proceed in stages, leaving the issue of voicing (order of voices from low to high) until later. Candidates can be rejected if they cause unschematized dissonances (Xs), or break some other constraint – e.g. we might reject parallel fifths, octaves, and unisons, to conform with the style. It turns out that we can construct viable fragments in which a cambiata, a neighbor, and a passing tone occur simultaneously, or in which a cambiata and two neighbors occur simultaneously (left as an exercise for the reader!). As far as we can tell, there's no "reason" that these don't occur in the corpus.

We can extend this game of finding unused potentials by taking the interval combinations of a structure as another parameter. For instance, four simultaneous passing notes occur about 40 times in the combined corpus, but most of the time the passing note "chord" is just a minor or major third, with pairs of passing notes up and down through each note of the third. There is one instance where a minor triad is constituted (one passing note is preceded by a dissonant third quarter). The major triad occurs several times in the triple-passing-tone structure; it appears to be an unused possibility in the quadruple.

The possibility for combinatorial explorations are vast. For instance, there are more than 70 different *sonorities* (pitch-class sets sounding at some moment) in Palestrina, while only 7 of them need not involve dissonance. The rest are constructed precisely in the manner we have just been describing, with combinations of dissonance treatments.

Equally great are the opportunities for musicologists to study specific usages in their musical, textual, and historical contexts; the computational means to find and annotate sets of occurrences will surely facilitate this process.

The general methodology used here can be extended to other corpora, and to other aspects of musical practice. The computational study of musical corpora through schematization, structure-building, automatic annotation, and generative extrapolation will bring a new scope and precision to our understanding of musical practice and potential.

[9] We can see this because these two instances have different polyphonic profiles: if they both resolve at the same time they're in rhythmic monophony with one another, whereas if one resolves first, one suspension is durationally contained within the other.

[10] I.e. a suspension with duration of a quarter.

[11] 261 movements of Victoria vs. 705 of Palestrina

[12] In fact the *whole* space of dissonance structures under this model may be small enough to be feasibly enumerable. If so, how does this fact relate to our surmise that *for Palestrina and Victoria*, the space seems to be "composed" rather than enumerated? This is a question for the practice and philosophy of the nascent discipline of constructive musicology, or the study of corpora through computational extension.

7. REFERENCES

[1] Fux, J.J.; Mann A. (trans. & ed.): *The Study of Counterpoint from J.J. Fux's Gradus Ad Parnassum.* Norton & Co. (1971/1965/1725)

[2] Jeppesen, K.: *Counterpoint: The Polyphonic Vocal Style of the Sixteenth Century.* Dover Publications (1992/1939/1931)

[3] Jeppesen, K.: *The Style of Palestrina and the Dissonance* Dover Publications (1970/1946)

[Sch.1999] Schubert, P.: *Modal Counterpoint, Renaissance Style* Oxford University Press (1999)