A FLEXIBLE APPROACH TO AUTOMATED HARMONIC ANALYSIS: MULTIPLE ANNOTATIONS OF CHORALES BY BACH AND PRÆTORIUS

Nathaniel Condit-Schultz
McGill University
nathaniel.condit-schultz@mcgill.ca

Yaolong Ju
McGill University
yaolong.ju@mail.mcgill.ca

Ichiro Fujinaga
McGill University
ichiro.fujinaga@mcgill.ca

ABSTRACT

Despite being a core component of Western music theory, *harmonic analysis* remains a subjective endeavor, resistant automation. This subjectivity arises from disagreements regarding, among other things, the interpretation of contrapuntal figures, the set of "legal" harmonies, and how harmony relates to more abstract features like tonal function. In this paper, we provide a formal specification of harmonic analysis. We then present a novel approach to computational harmonic analysis: rather than computing harmonic analyses based on one specific set of rules, we compute all possible analyses which satisfy only basic, uncontroversial constraints. These myriad interpretations can later be filtered to extract preferred analyses; for instance, to forbid 7th chords or to prefer analyses with fewer non-chord tones. We apply this approach to two concrete musical datasets: existing encodings of 371 chorales by J.S. Bach and new encodings of 200 chorales by M. Prætorius. Through an online API users can filter and download numerous harmonic interpretations of these 571 chorales. This dataset will serve as a useful resource in the study of harmonic/functional progression, voice-leading, and the relationship between melody and harmony, and as a stepping stone towards automated harmonic analysis of more complex music.

1. INTRODUCTION

Broadly, *harmony* refers to the simultaneous sounding of multiple pitches [22]. However, harmonic theory involves far more than just pitch collections. Rather, harmonic theory describes an abstract syntactic structure in Western tonal music, hierarchically removed from the literal pitches of the musical "surface" [22]. Though harmonic theory is a foundational component of basic music theory, the details of the theory are vague, and deceptively complex [5]. Harmony intertwines low-level sensory distinctions (consonance vs dissonance), short-term musical constructs (counterpoint, voice-leading), and abstract long-range musical structures (function, form, tonality, etc.), and thus plays a central role in musical experience. Given this complexity, it is no surprise that actual harmonic analysis is highly subjective, and thwarts any attempt to systematize or automate it. This paper attempts to clarify the dimensions of harmonic analysis, identifying the import points of disagreement and ambiguity in harmonic theory. We then present a novel approach to automated harmonic analysis, which allows us to generate a variety of consistent harmonic annotations based on a various assumptions and preferences.

Attribution: Nathaniel Condit-Schultz, Yaolong Ju, Ichiro Fujinaga. "A Flexible Approach to Automated Harmonic Analysis: Multiple Annotations of Chorales by Bach and Prætorius", 19th International Society for Music Information Retrieval Conference, Paris, France, 2018.

1.1 Theory and Terminology

To avoid confusion with the more general concept of "harmony," we use the term *sonority* to refer to pitch-class collections. The most basic sonority is the *dyad*—pairs of pitch classes which form *consonant* or *dissonant* intervals. [1] Larger sonorities can be seen as combinatorial compositions of dyads, as each new pitch class forms an interval with every other pitch class in the sonority. Harmonic theory generalizes about various dyad combinations, reducing a huge variety of possible interval combinations to a few categories. The central harmonic categories of Western music are the set of cardinal-three sonorities in which all intervals are consonant (*triads*) and the cardinal-four sonorities which include one dissonant interval (*7th chords*). Other sonorities—the preponderance of possibilities—are unclassified and considered nonsyntactic. Some genres (e.g., jazz, music theatre) employ larger sonorities (9th chords, 13th chords, etc.), which necessarily contain more dissonant intervals, as well as dissonant cardinal-three and cardinal-four sonorities (sus4, add9, etc.) [5], but even in these styles the vast majority of sonorities are unclassified.

Traditionally, dissonant harmonic intervals are only used in highly constrained melodic settings: Dissonant notes must "decorate" a neighboring consonant note, typically by moving to/from the consonance by step—a dissonance moving to a consonance by step is said to *resolve* to the consonance. Thus, a basic hierarchical distinction is introduced into music, as "decorative" dissonances are necessarily subservient to "structural" consonances. Traditional theory and pedagogy approaches larger musical textures by applying two-part concepts (parallelism, motion

---

[1] Here we only consider *generic* intervals, and thus *generic* dissonances. Generically, thirds, fifths, and sixths are consonant, though some specific versions of these intervals (e.g., diminished fifths, augmented thirds) are dissonant.

types, dissonance resolution) to all individual pairs. Unfortunately, larger textures introduce complexities which two-voice theory cannot address: decorative tones may appear in multiple voices at different times, at the same time, or even staggered such that one voice's decorative dissonance cooccurs with another's consonant resolution. As a result, the consonant harmonies which undergird musical syntax may never be explicitly sounded as sonorities. The concrete distinction between consonant and dissonant intervals gives way to a nebulous distinction between *chord-tones* which instantiate the local harmony and *non-chord* tones that decorate them [22]. This distinction, is the essential task of harmonic analysis. Traditional "roman numeral" harmonic analysis also requires some interpretation of higher-level tonal structures, including the global key and local modulations. Just as the melodic surface elaborates the underlying harmonic progressions, harmonic progressions elaborate more abstract *functional* (tonic, subdominant, dominant) progressions and prolongations, which in turn articulate the key or progressions between keys. This hierarchy, however, is not clear-cut or discrete: disentangling surface features from increasingly abstract structural progressions is difficult, and the procedure poorly defined.

1.2 Literature

Computational research into harmonic progression and function has been extensive [8, 12, 22–24, 27, 30]. Many researchers have sought to automate harmonic analysis, either using rule-based algorithms [9, 11, 21, 28, 29] or machine learning [13, 18, 19, 26]. Impressive performance has been achieved, though proper evaluation is somewhat difficult given that the "correct answer" is not clear cut. Even if interpretive leeway is allowed, algorithms inevitably struggle with even mildly idiosyncratic or exceptional passages—devising sufficiently complicated rules to cover all possibilities is impossible, and such passages are too rare to be learned by machine learning. Due to these difficulties, many researchers have relied instead on manual annotation by experts, who can make more nuanced decisions and adapt to never-before-seen situations [1,3,4,8,20]. However, though human analysts may create more accurate data, manual harmonic annotations—even by the same analyst—can be extremely inconsistent [14]. Given the subjectivity of harmonic analysis, the consistency of data annotation may actually more important than a vaguely-defined "accuracy" [6]—inconsistent answers to similar or identical musical patterns will inevitably hamper learning, whether human or machine.

To account for inconsistency and disagreement between theorists, many studies have employed multiple independent annotators [3, 4]. This approach is appropriate to the extent that analytical inconsistency is considered random noise. However, as we will explain, harmonic indeterminacy is not simply a matter of random error, but rather reflects fundamental disagreements concerning the nature, meaning, and purpose of harmonic analysis. Thus, annotation error is not (entirely) stochastic, but rather, is systematic. What's more, though multiple independent annotations give us some sense of the scope of disagreement between analysts, they do little to clarify the root causes of these disagreements. Our view is that is preferable to: A) precisely describe the subjective features of harmonic theory; B) study how different theoretical assumptions result in different analyses; and C) evaluate how well different assumptions/models explain patterns in music. The goal of our project is to facilitate these tasks.

1.3 Analytical ambiguity

Harmonic analysis is evidently a useful tool in the description of musical structure and musical experience, yet in practice, harmonic theory is underspecified with regards to many musical passages. Indeed, many prominent theories of music (e.g. Rameau, Riemann, Schenker) differ fundamentally in their approach to harmony. It is often possible to interpret the same passage in a number of ways. Furthermore, the informative distinctions conveyed by different interpretations is unclear. This ambiguity mainly regards four questions:

1. Which harmonies are "legal" structural harmonies? Are sevenths chords true harmonies, or are they always decorative?
2. How do we interpret sonorities that are subsets, supersets, or intersections of each other? Traditional harmonic categories like $\{\mathtt{V}, \mathtt{V}^7, \mathtt{vii}_o\}$ both share many pitch classes and have similar musical function—what, if any, useful information is conveyed by treating them as independent categories?
3. How do we interpret contrapuntally decorative notes which are consonant—i.e, can there be consonant non-chord tones? This issue is especially difficult when multiple voices engage in decorative motion at once, creating "passing chords."
4. Should harmonic analysis reflect only "surface" features (like dissonance resolutions), or should higher-level structures also play a role? For instance, should, large-scale parallelism inform analyses?—i.e., analyzing two parallel passages in a similar way even if their surface details differ?

Figure 1 illustrates a number of these issues in a concrete musical example. In Figure 1, the three notes colored red form dissonances and therefore *must* be interpreted as non-chord tones. Notes colored blue are consonant, but evince melodic contours similar to the dissonant notes. Throughout this paper, we refer to each new sonority formed whenever any voice articulates a new onset as a sonority "slice"—in Figure 1, slices are numbered above the grand staff.

The consonant passing tones in slices 2 and 8 are especially illustrative. If the passing tone in slice 2 is considered a chord tone, slices 1–2 form the harmonies $\mathtt{I} \rightarrow \mathtt{vi}_6$, both tonic function chords. If the passing tone in slice 8 is interpreted as a chord tone, the progression $\mathtt{ii} \rightarrow \mathtt{vii}_o^6$ results—a transition between two different tonal functions (subdominant and dominant). Given these functional differences, many analysts would mark slices 1–2 as a single `I` chord but slices 7–8 as $\texttt{ii}\rightarrow\texttt{vii}^{6}_{o}$. This is especially true since transitions from $\texttt{ii} \rightarrow \texttt{I}$ (slice 9) are considered abnormal, while transitions from $\texttt{vii}^{6}_{o} \rightarrow \texttt{I}$ are normative. Several slices illustrate the ambiguity regarding 7th chords: Passing tones in slices 6, 18, 20, 22 and 24 might each be interpreted as sevenths, or not. For instance, the G in slice 11 can be seen as the 7th of a $\texttt{ii}^{6}$ harmony, or as a suspension. In Bach's chorale music, chordal 7ths are nearly always treated like dissonances, begging the question: what is the difference between a "chord-tone 7th" and a "non-chord-tone 7th"?

2. CURRENT PROJECT

This paper describes a new approach to automated harmonic analysis, which remains agnostic regarding many of the specific interpretive complexities discussed so far. Rather, we base analyses on only a few, basic, uncontroversial constraints, allowing us to produce numerous interpretations of the same sonorities. Using this approach, we have generated a novel form of harmonic analysis dataset, including numerous harmonic annotations of chorales by Michael Prætorius (1571–1621) and Johann Sebastian Bach (1685–1750). This dataset can serve several useful functions:

1. Researchers can generate specific, *consistent* harmonic analyses, conforming to whatever analytical preferences/assumptions they prefer, for all music in the corpora. These analyses can be used like any other harmonic annotation data—i.e., to study harmonic progression and tonality in general.
2. The dataset includes a set of late-modal (Prætorius) and early-tonal (Bach) music, which are nonetheless largely similar in texture and style. This makes the dataset particularly useful for historical research [8].
3. Finally, by comparing analyses generated with different constraints, we can rigorously explore the ways in which different harmonic theories fit, or don't fit, real music.

Chorale music is invaluable for teaching and studying harmony, as it features consistent and highly constrained melodic/contrapuntal textures, with few non-chord tones. Bach's 371 chorales are mainstays of music theory pedagogy and have been the subject of much music information retrieval research [2, 7, 8, 16, 22–24, 27, 31]. Prætorius' 200 chorales are music of a somewhat similar texture, but have received relatively little attention. Several sets of expert analyses of Bach's chorales have been published, though only subsets of the chorales have annotations digitally aligned with symbolic music data. Other researchers have generated harmonic annotations—or analogous functional analyses—of the chorales computationally, and used these analyses in research, but have not published their annotations, nor describe them in detail.

3. METHODOLOGY

The approach of this project is to calculate all legal harmonic interpretations of a passage, and to only filter out specific interpretations at a later stage. Our approach is designed specifically for our dataset, and is thus rather "over fit" to chorale music, so it will not generalize well to other music. However, the basic concepts of our approach could be adapted to other tonal music.

Key to our entire endeavor is establishing "basic" constraints on harmonic interpretation. In true music theory form, we formulate these constraints as the following "rules." There are two types of rules: harmonic rules and melodic rules. Our harmonic rules are as follows:

1. Every sonority slice belongs to one and only one harmony.
2. Every new harmony must be followed by another new harmony on the next stronger metric position—i.e., harmonic rhythm cannot be syncopated. (Some Prætorius chorales contain exceptions to this rule, as the entire rhythmic texture is syncopated.)
3. Only triads (major, minor, diminished, or augmented) and 7th chords (dominant, major, minor, half-diminished, or fully-diminished) are considered legal harmonies. However, subsets of legal harmonies may also appear in music. Complete harmonies are preferred, but cardinal-three subsets of seventh chords (Root-3rd-7th or Root-5th-7th), dyadic subsets of triads (i.e., consonant intervals), and even unisons/octaves are permitted.

Given these definitions of harmony, we can then establish which notes do, or do not, belong to the local harmony. To be a non-chord tone, a note must satisfy the following melodic rules—any note that fails any of these rules *must* be a chord tone:

1. The antecedent and consequent note of each non-chord tone must be consonant (chord tones), excepting the special case of Rule 4g (below).
2. Non-chord tones cannot sustain across metric positions that are stronger than their own metric position.
3. Non-chord tones cannot sustain through changes of harmony. A note cannot start as a non-chord tone and then become a chord tone (though the opposite is possible, in the suspension).
4. Finally, all non-chord tones must match one of these traditional contrapuntal dissonance models:
    (a) *Passing tone*: approached and departed by step in the same direction.
    (b) *Neighbor tone*: approached and departed by steps in opposite directions; the antecedent and consequent are the same note.
    (c) *Suspension/Retardation*: approached by unison (or sustain); departed by step; stronger metric position than antecedent.

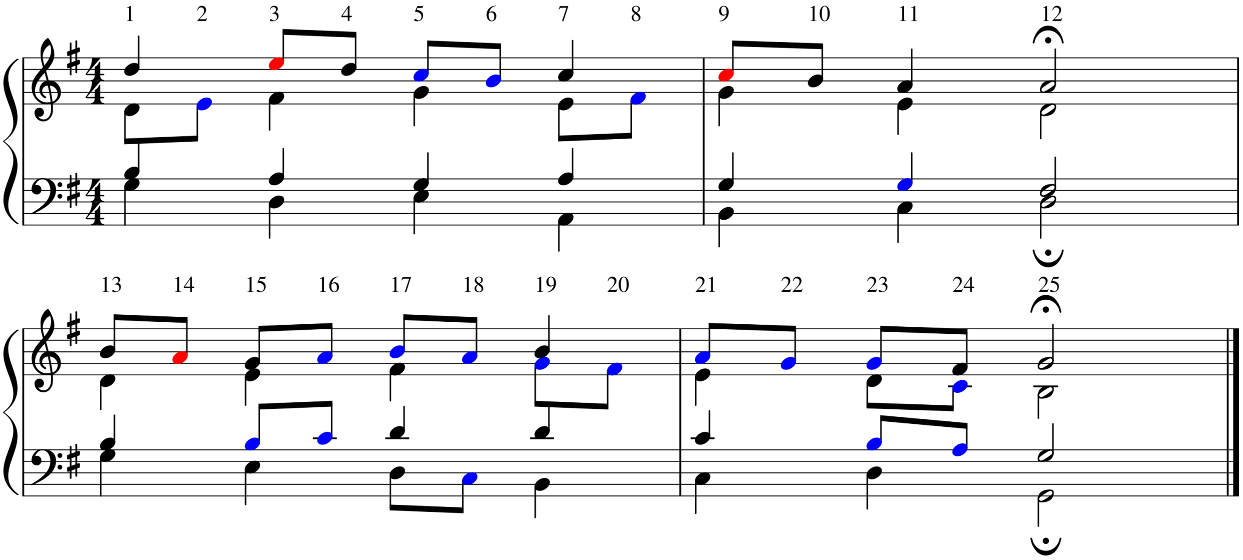

Figure 1. Illustration of "decorative" melodic idioms in a contrived example of four-part counterpoint. Slices (sonorities) are numbered above the staff. Notes colored red indicate dissonances. Notes colored blue indicate consonant notes which nonetheless articulate decorative melodic idioms, including passing tones (slices 2, 5, 8, 16, 18, 20, 21, 23, 24), neighbor tones (slices 6, 17, 18, 19), suspensions (slices 11, 23), a retardation (slice 15), and an anticipation (slice 22). Some of these interpretations are mutually exclusive, as a decorative tone cannot decorate another decorative tone. For instance, if the C in slice 5 is considered a passing tone, then the B in slice 6 must be a chord tone which resolves the passing tone.

(d) *Appoggiatura*: approached by leap; departed by step in opposite direction; stronger metric position that antecedent.

(e) *Escape tone*: approached by step; left by skip; weaker metric position than its antecedent.

(f) *Pedal tone*: approached by unison (or sustain); left by unison (or sustain).

(g) *Double passing*: two non-chord tones of the same duration, separated by step; approached and departed by step in the same direction; the first of the pair must occupy a weaker beat than its antecedent.

As in all dimensions of harmonic analysis, there is not universal agreement regarding the rules for non-chord tones. The rules set out here are an amalgam of the rules explicitly, or implicitly, described in typical music theory text books [15, 17], specialized (through some trial an error) for our chorale datasets.

3.1 Data parsing

Symbolic encodings of the Bach chorales were gathered from the KernScores repository (`kern.ccarh.org`), which is maintained by Stanford's Center for Computer Assisted Research in the Humanities. The music of 370 four-part chorales, and one five-part chorale [2], is encoded in the humdrum `**kern` representation (`www.humdrum.org`) [10]. Symbolic encodings of 200 chorales by Prætorius were recently digitized by members of McGill University SIMSSA project: Scores were initially scanned and interpreted by optical music recognition software before being corrected by a human annotator. This data was originally encoded in `musicXML` format, but was converted to `**kern` data for this project, so as to facilitate alignment with harmonic annotations. The Prætorius data includes 197 four-voice chorales and three five-voice chorales. In total, the dataset includes 571 chorales, consisting of 129,568 notes (+ 898 rests), which form 42,895 sonority slices.

`**kern` data was parsed using the Humdrum Toolkit [10], before being loaded into `R` [25] for additional parsing. The analysis workflow was also conducted in `R`. To make the analyses useful as comparisons across the two composers, (almost) the exact same parsing and analysis workflow are applied to each.

In addition to pitch and rhythm data, the Bach chorale data contains some phrasing information, in particular, fermatas. A phrase ending in a Bach chorale was identified whenever all four voices reach a fermata. [3] The Prætorius chorale data contains phrasing information, encoded as rests in all voices, and both datasets contain metric information. [4]

3.2 Workflow

Our process has a two-stage workflow. The first-stage is to divide the music exhaustively into contiguous groups

---

[2] This five-part chorale was excluded from the dataset available on Kernscores, but was encoded for the purposes of this study

[3] Several chorales had notational inconsistencies, wherein fermatas were not encoded on the inner voices. These inconsistencies were fixed manually.

[4] Though metric indications in Prætorius' era are not exactly conceptually equivalent to modern time signatures.

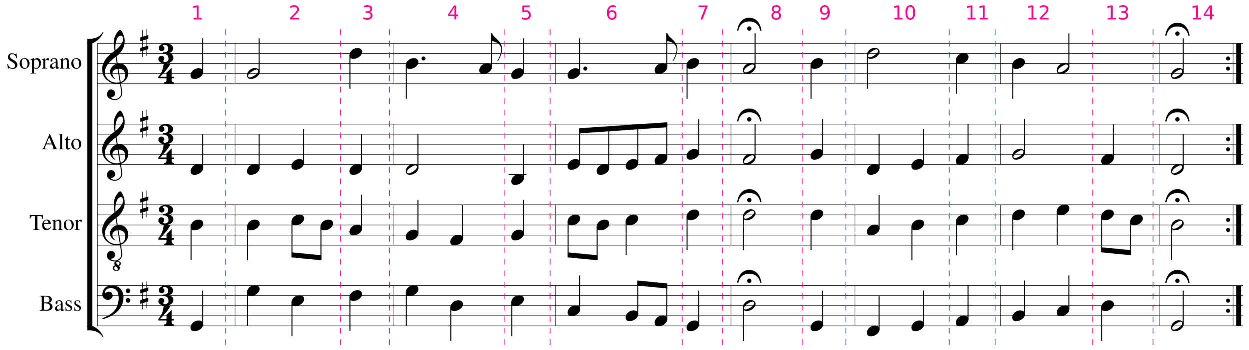

Figure 2. Illustration of contextual windows in Bach's Chorale 1, *Aus meines Herzens Grunde*. Slices between dashed red lines are analyzed as one window.

of successive slices: "contextual windows." The second-stage applies an analysis algorithm to each segment.

3.2.1 *Stage 1*

Many sonority slices can be analyzed in isolation. However, many more slices need context to by analyzed. Our approach is to parse the music into a single set of contiguous (non-overlapping) windows, identified using a simple, rule-based heuristic. A new contextual window begins anytime:

1. All voices attack on a strong beat.
2. All voices attack and one or more voices did not attack in the previous slice.
3. In an offbeat slice, more than two voices attack and one or more voices sustains into/past the next beat.
4. After a phrase boundary.

Figure 2 illustrates the contextual windows derived by this heuristic in the first seven measures of Bach's first chorale. The aim of this heuristic is to err on the side of larger segments: unnecessarily large windows can be broken down into separate harmonies at a later stage, but windows that are too small will not provide enough context to identify all legal interpretations, and in some cases may result in windows that are not parsable.

3.2.2 *Stage 2*

Once analytical windows are identified, we apply the following permutational algorithm to the slices in each window.

1. Identify all ways in which the window can be divided exhaustively into sub-segments while obeying harmonic-rhythm constraints (Harmonic Rule 2).
2. For each possible segmentation, identify all pitches that can legally be non-chord tones (Melodic Rules 3–4)—we call these *potential* non-chord tones.
3. Compute every combination of potential non-chord tones, allowing that some interpretations are mutually exclusive (detailed explanation below).
4. For every legal combination of potential non-chord tones, remove these non-chord tones and group the remaining chord tones into every possible sub-segment.
5. Discard interpretations which contain (any) illegal harmonies.
6. If any preferred harmonies are present, discard incomplete harmonies (Harmonic Rule 3).
7. If the same chord is identified in two successive slices, discard this interpretation (a different sub-segmentation is sure to have found the equivalent).
8. If a slice is identified as a dyad/unison, and the preceding or succeeding slice is a superset of that dyad/unison, the slice is subsumed into the superset.

Figure 3 illustrates the application of this algorithm to the sixth window in the chorale shown in Figure 2. The four slices in this window can be legally divided in six segmentations (Step 1), shown below the staff. Eleven of the twelve notes in the window are potential non-chord tones (labeled and enumerated in Figure 3). The algorithm tests various permutations of these potential non-chord tones (algorithm Steps 3–5) as so: First, assume potential non-chord #1 *is* a non-chord tone and all other notes are chord tones: under this assumption, segmentation `1...` forms the illegal sonority $\{A,B,C,D,E,F\#\}$; segmentation `1.2.` forms the illegal sonorities $\{B,C,D,E\}$ and $\{A,B,C,E,F\#\}$; segmentation `1..2` forms the illegal sonority $\{B,C,D,E\}$ and the legal sonority $\{A,C,F\#\}$; etc. Repeat this procedure for every other potential non-chord tone, every pair of non-chord tones, every triplet of non-chord tones, etc., skipping combinations which are mutually exclusive—i.e., if #2 is an appoggiatura #4 must be a chord tone (Rule 1). Testing all non-chord tone permutations across all six segmentations reveals eleven non-redundant (Steps 6–8) interpretations with legal chords in all segments (Step 5).[5] Of these eleven, we can "filter out" interpretations involving 7th chords, leaving the three triadic analyses shown in Figure 3.

---

[5] Our actual algorithm incorporates a few additional optimizations to limit the number of permutations which must be tested. The most important involves pitch classes: within a given harmonic segment all instances of a single pitch class must be either non-chord tones or chord tones. For instance, it would be meaningless to treat #9 as a passing tone but treat #11 as a chord tone. Similarly, #7 (a C) can never actually be a passing tone, since the C in the bass is always a chord tone.

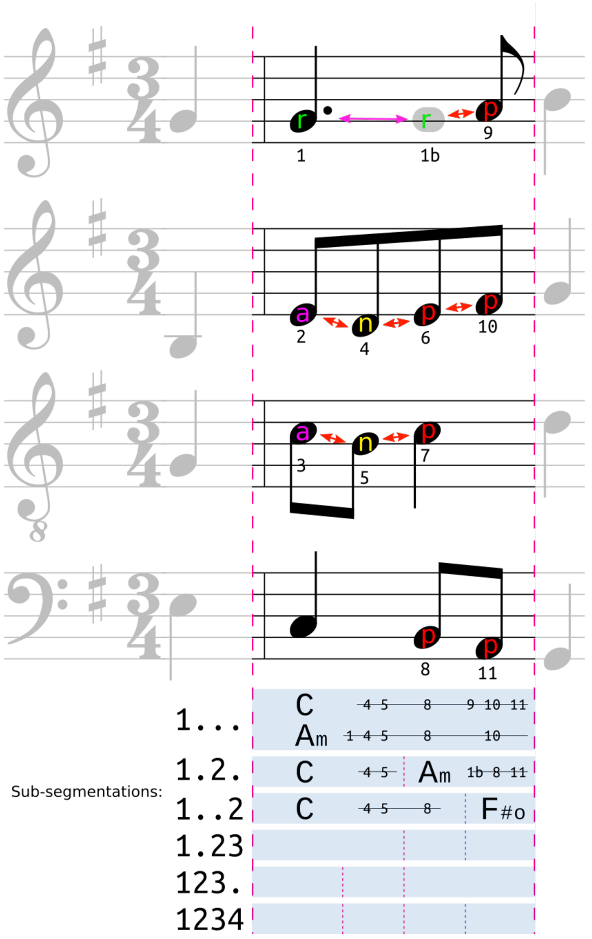

Figure 3. Illustration of the permutational analysis of a single contextual window (window 6 from Figure 2). Each note in the window is annotated as a potential non-chord tone, marked `p` for passing tone, `n` for neighbor tone, `r` for retardation, or `a` for appoggiatura—mutually exclusive potentials are annotated with arrows. The single unlabeled C *must* be a chord tone, as it does not match any contrapuntal dissonance model (Melodic Rules 4). Below the staff, the six possible rhythmic segmentations of the window are shown. The four possible purely-triadic interpretations of the window are show; the notes which are interpreted as non-chord tones are identified (by number) beside each analysis.

3.3 Edge cases

Chorale music is valued pedagogically for its simplicity and consistency. Nonetheless, a handful of chorales contain unusual features which complicate the batch analysis of the corpora. Notable examples in the Bach chorales include: an unusual call and response between the soprano and the rest of the voices in Chorale 43; dissonant notes which resolve across phrase boundaries (i.e., through a fermata) in Chorales 127, 202, and 234; and suspensions which resolve indirectly in Chorales 5 and 199. A number of Prætorius chorales also contain subsections in which a subset of voices sing while the others rest, confounding our windowing heuristic. Solutions to these special cases, and a handful others, were hard-coded into the workflow.

4. API

The data is hosted at `github.com/DDMAL/Flexible_harmonic_chorale_annotations`. The harmonic permutation data is stored in a `rData` file. Users may filter out specific harmonic analyses using an online GUI, and download them as a zipped collection of text files encoded in the Humdrum Syntax. Each file contains the `**kern` representation of a chorale aligned with one or more harmonic analyses in a `**harm` representation. Interpretations can be filtered by the following criteria:

- Type of harmonies.
- Number of harmonies (per beat/per window).
- Types of non-chord tones.
- Number of non-chord tones (per slice/per window).

For example, one could extract analyses which forbid augmented triads, appoggiaturas, and ♪ harmonic rhythms. Users may also download the raw data and associated `R` scripts for local use or customization.

5. CONCLUSION

The empirical and computational study of harmony is essential to furthering our understanding of musical structure and perception. However, this research must remain cognisant of the subtle complexities and controversies of harmonic theory if it is to be fruitful. We have presented a novel approach to automated harmonic analysis which is not limited to one specific set of theoretical assumptions, allowing for just such subtleties to be explored systematically. We have also described a new dataset generated via this method. We hope that this dataset will facilitate research into tonality and harmonic progression, especially changes in harmonic practice between the early 1600s and the mid 1700s. However, our grander purpose is to facilitate critical, data-driven, interrogation of harmonic theory in general.

6. REFERENCES

[1] John Ashley Burgoyne, John Wild, and Ichiro Fujinaga. An Expert Ground Trught Set for Audio Chord Recognition and Music Analysis. In A Klapuri and C. Leider, editors, *Proceedings of the 12th International Society for Music Information Retrieval (ISMIR) Conference*, pages 633–638, Miami, FL, 2011.

[2] Trevor de Clercq. A Model for Scale-Degree Reinterpretation: Melodic Structure, Modulation, and Cadence Choice in the Chorale Harmonizations of J. S. Bach. *Empirical Musicology Review*, 10:188–206, 05 2015.

[3] Trevor de Clercq and David Temperley. A Corpus Analysis of Rock Harmony. *Popular Music*, 30(01):47–70, January 2011.

[4] Johanna Devaney, Claire Arthur, Nathaniel Condit-Schultz, and Kirsten Nisula. Theme and Variation Encodings with Roman Numerals: A New Data Set for Symbolic Music Analysis. In Meinard Müller and Frans Wiering, editors, *Proceedings of the 16th International Society for Music Information Retrieval (ISMIR) Conference*, pages 728–734, Malaga, Spain, 2015.

[5] Christopher Doll. Definitions of 'Chord' in the Teaching of Tonal Harmony. *Dutch Journal of Music Theory*, 18(2):91–106, 2013.

[6] Mark Granroth-Wilding and Mark Steedman. A Robust Parser-Interpreter for Jazz Chord Sequences. *Journal of New Music Research*, 43(4):355–374, 2014.

[7] Gaëtan Hadjeres and François Pachet. DeepBach: a Steerable Model for Bach Chorales Generation. *CoRR*, 2016.

[8] Thomas Hedges and Martin Rohrmeier. Exploring Rameau and Beyond: A Corpus Study of Root Progression Theories. In *Mathematics and Computation in Music: Third International Conference, MCM 2011, Paris, France, June 15-17, 2011. Proceedings*, pages 334–337, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[9] Tim Hoffman and William P Birmingham. A Constraint Satisfaction Approach to Tonal Harmonic Analysis. *Electrical Engineering and Computer Science Department. Technical Report CSE-TR-397-99. University of Michigan*, 2000.

[10] David Huron. *Music Research Using Humdrum: A User's Guide*. Stanford, California: Center for Computer Assisted Research in the Humanities, 1999.

[11] Plácido R. Illescas, David Rizo, and José M. Iñesta. Harmonic, Melodic, and Functional Automatic Analysis. In *Proceedings of the International Computer Music Conference*, pages 165–168, 2007.

[12] Nori Jacoby, Naftali Tishby, and Dmitri Tymoczko. An Information Theoretic Approach to Chord Categorization and Functional Harmony. *Journal of New Music Research*, 44(3):219–244, 2015.

[13] Yaolong Ju, Nathaniel Condit-Schultz, and Ichiro Fujinaga. Non-chord Tone Identification Using Deep Neural Networks. In *Proceedings of the Fourth International Workshop on Digital Libraries for Musicology*, pages 13–16. ACM, 2017.

[14] Hendrik Vincent Koops, W. Bas de Hass, John Ashley Burgoyne, and Jeroen Bransen. Harmonic Subjectivity in Popular Music. Technical report, Department of Information and Computing Sciences, Utrecht University, 2017.

[15] Stefan Kostka and Dorothy Payne. *Tonal Harmony: With an Introduction to Twentieth-Century Music*. McGraw Hill, 2004.

[16] Pedro Kröger, Alexandre Passos, Marcos Sampaio, and Givaldo De Cidra. Rameau: a System for Automatic Harmonic Analysis. In *Proceedings of International Computer Music Conference*, pages 273–281, 2008.

[17] Steven G. Laitz. *The Complete Musician: an Integrated Approach to Tonal Theory, Analysis, and Listening*. Oxford University Press, 3rd edition, 2012.

[18] Kristen Masada and Razvan Bunescu. Chord Recognition in Symbolic Music Using Semi-Markov Conditional Random Fields. In *Proceedings of the 18th International Society for Music Information Retrieval Conference*, pages 272–278, 2017.

[19] Lesley Mearns. *The Computational Analysis of Harmony in Western Art Music.* PhD thesis, Queen Mary University of London, 2013.

[20] Néstor Nápoles López. Automatic Harmonic Analysis of Classical String Quartets from Symbolic Score. Master's thesis, Universitat Pompeu Fabra, 2017.

[21] Brian Pardo and William P. Birmingham. The Chordal Analysis of Tonal Music. In *The University of Michigan, Department of Electrical Engineering and Computer Science Technical Report CSE-TR-439-01*, 2001.

[22] Ian Quinn. Are Pitch-Class Profiles Really Key for Key? *Gesellschaft für Musiktheorie*, 7(2):151–163, 2010.

[23] Ian Quinn and Panayotis Mavromatis. Voice-Leading Prototypes and Harmonic Function in Two Chorale Corpora. In Carlos Agon, Moreno Andreatta, Gérard Assayag, Emmanuel Amiot, Jean Bresson, and John Mandereau, editors, *Mathematics and Computation in Music*, pages 230–240, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[24] Ian Quinn and Christopher Wm. White. Expanding Notions of Harmonic Function Through a Corpus Analysis of the Bach Chorales. In *Annual Conference of the Society for Music Theory*, Charlotte, NC, 2013.

[25] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013.

[26] Christopher Raphael and Joshua Stoddard. Functional Harmonic Analysis Using Probabilistic Models. *Computer Music Journal*, 28(3):45–52, 2004.

[27] Martin Rohrmeier and Ian Cross. Statistical Properties of Tonal Harmony in Bach's Chorales. In *Proceedings of the 10th International Conference on Music Perception and Cognition*, 2008.

[28] Craig Stuart Sapp. Computational Chord-Root Identification in Symbolic Musical Data: Rationale, Methods, and Applications. *Computing in Musicology*, 15:99–119, 2007.

[29] David Temperley and Daniel Sleator. Modeling Meter and Harmony: A Preference-rule Approach. *Computer Music Journal*, 23(1):10–27, 1999.

[30] Christopher Wm. White. Changing Styles, Changing Corpora, Changing Models. *Music Perception*, 31(3):244–253, 2014.

[31] Matthew Woolhouse. Probability and Style in the Chorales of J. S. Bach. *Empirical Musicology Review*, 10(3):207, 2015.