ZGMTH Zeitschrift der
Gesellschaft für Musiktheorie

Quinn, Ian (2010): Are Pitch-Class Profiles Really “Key for Key”? ZGMTH 7/2, 151–163. https://doi.org/10.31751/513

veröffentlicht / first published: 05/07/2010
zuletzt geändert / last updated: 31/08/2010

Are Pitch-Class Profiles Really "Key for Key"?[1]

Ian Quinn

Most current approaches to key-finding, either from symbolic data such as MIDI or from digital audio data, rely on pitch-class profiles. Our alternative approach is based on two ideas: first, that chord progressions, understood rather loosely as pairs of neighboring harmonic states demarcated by note onsets, are sufficient as windows for key-finding, at least in the chorale context; and second, that the encapsulated identity of a chord progression (modulo pitch-class transposition and revoicing) is sufficient – that is, that reduction of progressions to pitch-class distributions is not necessary for key-finding. The system has no access to explicit information about a chord progression other than its transpositional distribution in the training corpus, yet it is able to reach an almost stunning degree of subtlety in its harmonic analysis of chorales it's never heard before. This suggests that reductionist approaches to tonality may be off the mark, or at least that pitch-class reductionism might not be necessary for a principled account of key.

Most current approaches to key-finding, either from symbolic data such as MIDI or from digital audio data, use some form of the following procedure developed by Krumhansl and Schmuckler (Krumhansl 1990):

- Select a window of music to be analyzed;
- Determine the distribution of pitch classes within the window;
- Use the pitch-class distribution to determine the most likely key.

The devil, as usual, is in the details. Should window size be defined in terms of chronological time, notational (metrical) time, or number of note onsets? How big should the window be? How does key-finding for a given window take into account results for previous windows? How should pitch-class distributions be weighted? How is a key determination made from a given pitch-class distribution? The consideration of these and related questions has dominated the key-finding literature for some time.[2]

This paper proposes a novel approach to key-finding that is not based on pitch-class distributions, at least not in any recognizable sense. Tailored to a genre (the four-part chorale) favored first by early-modern theologians and then by late-modern music-theo-

[1] I am grateful to Ian Bates and James Hepokoski for helpful commentary on an earlier version of this paper.

[2] The title of this article alludes to Temperley 1999, one of the first articles to question Krumhansl's methodology. For a more recent example of a discussion of methodological details see Chuan and Chew 2005.

ry pedagogues, this approach is based on two ideas: first, that chord progressions, understood rather loosely as pairs of neighboring harmonic states demarcated by note onsets, are sufficient as windows for key-finding, at least in the chorale context; and second, that the encapsulated identity of a chord progression (modulo pitch-class transposition and revoicing) is sufficient to serve as a basis for key-finding – that is, that knowledge of pitch-class distributions is not necessary for key-finding.

1. First data structure: the chord.

The reader is hereby forewarned that the word chord will be used in a rather nonstandard sense in this paper. In its ordinary-language sense, the word 'chord' is often used to refer to a harmonic entity with features that can be used to compare it with other chords: a root, a quality (e. G., major, minor, dim7, sus4), and possibly an inversion. In its ordinary-language sense, not any collection of pitch-classes can be a chord – only those that can be generated by stacking intervals of a third (relative to some diatonic scale). 'A' chord, in this context, is typically a theoretical entity abstracted away from the musical surface: although every note in a composition is said to fall under the influence of a chord, non-chord tones can appear, usually with the expectation that they will resolve in some style-dependent way.

1.1. Definition

Our sense of the word 'chord' will differ substantially from its ordinary pedagogical usage, referring to a snapshot of all pitch-classes sounding at any given moment.[3] (In the context of MIDI data, which constitute the input to the key-finding algorithm discussed here, the chord corresponding to a given time is determined by the set of note-on messages active at that time; things are slightly more complicated in the context of unperformed notation, and significantly more complicated in the case of audio data.) There is therefore no distinction between chord tones and non-chord tones: a chord is no more or less than the sum of the tones heard. Nor is there a restriction on what can constitute a chord; the model knows no distinction between consonant and dissonant, diatonic and chromatic, tertian and quartal, or any other set of theoretical terms.

Related to the model's lack of distinction between chord tones and non-chord tones is its radically localized conception of what constitutes a chord: simply put, every time a new note sounds, a new chord is identified. The nature of the four-part chorale virtually guarantees that exactly four notes are sounding at any given moment; with few exceptions, each MIDI note-off message is immediately followed by a notionally simultaneous note-on message. Clearly it is this property of chorales that gives us the latitude to define chord in a manner that is at once so broad and so black-and-white. We will set aside

[3] David Temperley uses a similar structure (the pitch-class set) in his Bayesian polyphonic key-finding algorithm (Temperley 2002, 207), but there are two important differences. First, Temperley's model uses a window rather than an instantaneous time-slice. Second, Temperley analyzes the pitch-class set into its constituent pitch-classes, whereas our model encapsulates chords.

for now the question of the extent to which this definition can be carried over to other genres and other (non-MIDI) encodings.

Having said what can and should be called a chord, it remains for us to specify how chords should be identified. A chord will be characterized with reference to the bass voice, defined in this repertoire with respect to a particular contrapuntal part that almost always carries the lowest sounding note. The upper parts are identified in terms of their intervallic relationship to the bass, modulo octave equivalence and permutation of the upper parts. In other words, each chord's essential properties are roughly the same as those embodied in figured-bass notation: no attention is paid to how notes above the bass are assigned to particular voices or registers. (The primary difference between figured-bass practice and the present model is that we reckon intervals above the bass in semitones rather than diatonic scale-steps.) We may pay attention to doublings – the number of copies of each pitch-class in each chord – or we may ignore this information.

Once the upper-voice intervallic structure of the chord is determined, we encapsulate this information via a human-readable label. From the key-finding algorithm's point of view, this label is essentially arbitrary and unparseable – it is not merely a pitch-class distribution profile in disguise. Example 1 illustrates the identification of chords in a Bach chorale.

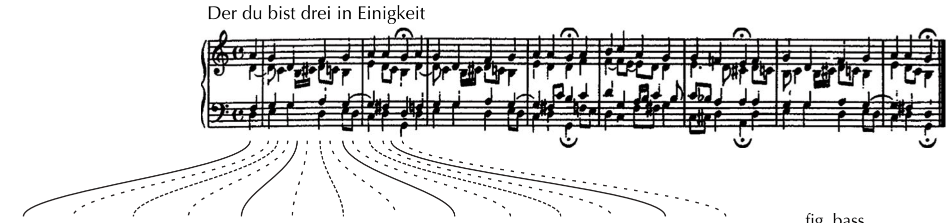

0,3,7	3,3,10	3,3,8	0,5,7	0,4,7	0,6,7	0,3,7	3,7,10	3,3,7	5,5,9	4,7,9	0,4,7	4,7,10	0,4,7	← fig. bass (with doubling)
7	9	9	0	0	0	7	7	9	7	5	7	7	0	← bass line
3,7	3,10	3,8	5,7	4,7	6,7	3,7	3,7,10	3,7	5,9	4,7,9	4,7	4,7,10	4,7	← fig. bass (without doubling)
mi.53	mi.73	ma.63	ss.54	ma.53	??.67	mi.53	mi.7	mi.53	ma.64	mi.65	ma.53	do.7	ma.53	← fig. bass, human readable (without doubling)

Example 1

The second row from the bottom shows the semitonal figured bass, and the bottom row shows the human-readable labels. Each label begins with a two-letter chord indicating the quality of the chord (see Table 1), followed by a dot and a series of numbers that specify the figured bass.[4]

[4] For ease of reading, these numbers are given as diatonic intervals when the chord quality is identifiable; for unknown or unspecifiable chord types the numbers indicate the semitone size of intervals above the bass.

letters	quality			numbers
	triad	*seventh-chord*	*other*	
ma	major	major	—	diatonic figured bass
mi	minor	minor	—	〃 〃 〃
di	diminished	diminished	—	〃 〃 〃
au	augmented	—	—	〃 〃 〃
do	—	dominant	—	〃 〃 〃
hd	—	half-diminished	—	〃 〃 〃
ss	—	—	suspension-derived	〃 〃 〃
pp	—	—	passing-derived	〃 〃 〃
??	—	—	other	chromatic figured bass

Table 1

1.2. Distribution of chords in the Bach corpus.

Of the 371 chorales in the Riemenschneider edition, twenty-two either duplicated another harmonization or contained phrases repeated more than twice; these were excluded from the study on the basis that including them might give a distorted picture of the distribution of chords and chord progressions in the corpus. The remaining 349 chorales, which constitute what we will refer to here as the 'Bach corpus,' contain 33,978 chords representing 167 distinct chord types. Table 2 collates these types into a familiar hierarchy of categories – chord type, chord quality, and inversion[5] – plus some catch-all categories. The 27 types corresponding to the familiar four triad qualities and five seventh-chord qualities account for slightly more than three-quarters of all the chords. Another 8 % or so of chords could be understood as representing 10 tertian types such as incomplete triads or seventh chords, chords with added ninths, and so on.[6] The remaining 15 % of chords represent 130 nontertian chord types resulting from what are ordinarily called 'non-chord tones,' though we reiterate that this model does not distinguish between chord tones and non-chord tones.

2. Second data structure: the chord progression

Just as we adopted a broad view of what is considered a chord, let us define chord progressions without worrying too much about what is usually meant by the term.

[5] Because the system assumes enharmonic equivalence, no distinction can be made between inversions of augmented triads or diminished seventh chords.

[6] It might be misleading to include ninth chords in the 'tertian' category, since these are almost all 9–8 suspensions, and the chords resulting from other kinds of suspensions (2–3, 4–3, and 7–6) are included in the 'nontertian' category.

root position	8054	67.5%		
first inversion	3451	28.9%		
second inversion	423	3.5%		
major triads	11928	100.0%	61.4%	
root position	3875	65.5%		
first inversion	1659	28.0%		
second inversion	386	6.5%		
minor triads	5920	100.0%	30.5%	
root position	206	14.1%		
first inversion	1099	75.4%		
second inversion	153	10.5%		
diminished triads	1458	100.0%	7.5%	
augmented triads	116		0.6%	
TOTAL TRIADS	19422		100.0%	57.2%
root position	1299	42.6%		
first inversion	995	32.6%		
second inversion	168	5.5%		
third inversion	589	19.3%		
dominant sevenths	3051	100.0%	44.1%	
root position	501	29.4%		
first inversion	526	30.9%		
second inversion	261	15.3%		
third inversion	417	24.5%		
minor sevenths	1705	100.0%	24.7%	
root position	381	40.4%		
first inversion	87	9.2%		
second inversion	135	14.3%		
third inversion	339	36.0%		
major sevenths	942	100.0%	13.6%	
root position	174	20.2%		
first inversion	339	39.4%		
second inversion	227	26.4%		
third inversion	120	14.0%		
half-diminished sevenths	860	100.0%	12.4%	
diminished sevenths	358		5.2%	
TOTAL SEVENTHS	6916		100.0%	20.4%
nontertian	5048	66.1%		
tertian	2592	33.9%		
TOTAL OTHER	7640	100.0%		22.5%

Table 2

2.1. Definition

We define a chord progression as a pair of chords (each characterized as a collection of intervals above the bass) plus the ordered pitch-class interval between the two bass notes, measured in semitones. For purposes of the key-finding algorithm, the identity of a chord progression is also encapsulated and labeled with an arbitrary but human-readable label. The statistical information the system uses to estimate key centers, in other words, does not include any specific information about what chords constitute a progression or what the intervals the bass line traverses. The label takes a form like this:

do.42–(10)–mi.63

which refers to a dominant seventh chord with its seventh in the bass, resolving normally to a minor triad in first inversion. The number 10 indicates that the bass moves down by step (–2 semitones, modulo 12).

2.2. Distribution of chord progressions in the Bach corpus.

In the 349 Bach chorales used in study (see section 1.2), there are 33,630 chord progressions representing 3353 distinct chord-progression types. The relative frequency of chord-progression types seems to obey Zipf's law; that is, a ranking of all progression types by frequency will show a power-law relationship between rank order and frequency. For Zipf's-law distributions, log frequency is a linear function of log rank, with a slope close to –1 (in this case, about –1.23). This sort of distribution is associated with a number of social and linguistic phenomena, including most significantly the frequencies of words in natural-language corpora. Curiously, the distribution of progressions follows Zipf's law much more closely than the distribution of chord types, suggesting that the best way to make analogies about 'tonal grammar' in this repertoire might be to think of progressions, not chords, as playing the role that words play in language. (Chords might be productively thought of as analogous to morphemes.)

Table 3 lists progression types occurring more than 200 times in the Bach corpus and provides information about the scale-degree distribution of each progression.

For each instance of a progression in the corpus, we determine the scale-degree identity of the first bass note relative to the chorale melody's final pitch class (which is assumed to be the tonic); the table tallies the percentage of instances of each progression that occur on any given scale degree. Consider, for example, the fourth-ranking progression (n = 544),

do.7–(5)–ma.53,

consisting of a root-position dominant seventh chord followed by a root-position major triad, with the bass moving down by fifth. In over half of the instances of this progression, the first bass note is scale degree 5 relative to the global tonic. In most of the remaining instances, the first bass note is a whole step above or below the tonic, which we would understand as scale degree 5 of a tonicized V (dominant) or III (relative major of a minor key).

				probability of beginning on scale degree												
rank	*n*	irrev	progression	−2	−6	−3	−7	4	1	5	2	+6	+3	+7	+4	k
1	886	0.98	ma.53-(0)-do.7	0	1	2	14	2	8	53	12	1	2	0	0	5.8
2	815	0.5	ma.53-(0)-ma.53	0	3	14	6	3	40	28	1	0	0	0	0	0.3
3	806	0.75	ma.53-(5)-ma.53	0	2	7	13	5	17	43	9	0	0	0	0	3.6
4	544	1	do.7-(5)-ma.53	0	0	1	15	3	4	57	16	0	0	0	0	5.3
5	505	0.91	ss.54-(0)-ma.53	0	4	4	12	6	13	42	8	2	2	0	0	5.0
6	364	0.62	ma.63-(1)-ma.53	0	0	0	0	0	2	7	19	8	17	36	6	1.1
7	331	0.5	mi.53-(0)-mi.53	0	0	0	2	8	59	7	3	13	2	0	0	6.0
8	315	0.93	ma.53-(10)-do.42	0	0	4	14	6	14	49	5	0	1	0	0	5.3
9	309	0.94	do.65-(1)-ma.53	2	0	0	0	0	1	6	13	10	10	29	24	2.5
10	300	0.99	mi.53-(10)-mi.42	0	0	0	2	9	42	7	6	26	2	0	0	1.2
11	300	0.66	ma.53-(5)-mi.53	0	0	0	0	3	22	51	2	5	11	0	0	3.4
12	281	0.97	ma.53-(11)-ma.42	4	10	12	4	16	45	3	0	0	0	0	0	4.3
13	272	0.97	do.42-(11)-ma.63	5	16	6	11	47	9	0	0	0	0	0	1	6.2
14	267	0.25	ma.53-(7)-ma.53	2	10	9	2	17	53	2	0	0	0	0	0	4.8
15	258	0.87	di.63-(10)-ma.53	0	0	1	7	15	12	5	31	22	0	0	0	-0.7
16	234	0.75	mi.53-(0)-di.63	0	0	1	6	23	10	4	34	15	2	0	0	-0.2
17	224	0.38	ma.53-(11)-ma.63	2	8	8	5	15	52	4	0	0	0	0	0	5.2
18	212	0.92	ma.63-(0)-do.65	4	9	1	0	0	2	4	6	11	14	32	10	2.9
19	211	0.91	ma.53-(0)-ma.7	2	8	13	4	38	28	1	0	0	0	0	0	-0.0
20	204	0.95	do.65-(1)-mi.53	12	18	1	0	0	0	0	0	4	18	29	12	0.9
21	200	0.99	mi.65-(2)-ma.53	2	20	3	2	48	19	0	0	0	0	0	0	3.8

Table 3

Because tonicizations and modulations are quite common in the Bach chorales, it is not typical for a progression to show a frequency much higher than 50 % for any given scale degree. Nor do many progressions have such sharply peaked scale-degree distributions as the one we have just seen. Consider two variants of the major V^7–I cadence just discussed. Without the seventh (see the third-ranked progression, $n = 806$), the descending-fifth progression between two major triads is still associated with scale degree 5, but less so; also, the progression tends to happen more frequently starting from scale degree 1 (tonic to subdominant rather than dominant to tonic). The ninth-ranked progression ($n = 309$), a variant of the V^7–I progression with the seventh chord in first inversion, has a much 'flatter' distribution than either of the other two, no doubt thanks to the frequent appearance of this progression in tonicizations and modulations instead of final cadences.

3. The key-finding algorithm

The present key-finding algorithm takes advantage of the differential distribution of progression types over scale degrees, leveraging this information to make accurate estimates of key center by means of a statistical-learning model. Since the model needs to be trained on real-world data, we avoid overfitting the model (i. e., begging the question) by dividing the Bach corpus a priori into a training set (the first 314 chorales, or 90% of the corpus, ordered by Riemenschneider number) and a test set (the last 35 chorales, or 10% of the corpus) which we can use to judge the model's performance.

3.1. Training phase.

We train the model through the same procedure used to generate Table 3. Chorales are fed into the model one at a time. The system assumes that the final pitch class of the chorale melody is the global key center of the chorale. For each progression in each chorale, the model determines the scale-degree identity of the progression's first bass note, based on the interval from the chorale's global tonic and the bass note in question, and collates this information in a table. Each row of the table corresponds to a progression type, and each column corresponds to one of the twelve possible chromatic scale degrees (up to enharmonic equivalence). For every observation of a progression, the system increments the value of one cell in the table. Once the system has processed all the chorales in the training set, tallying each instance of each progression it encounters, it converts the observations into probabilities by dividing each cell by the row total. The table then resembles an expanded version of Table 3. Each row represents a probability distribution that allows the system to answer the following question: Given a particular instance of this progression in a novel chorale from the same corpus as the training set, and all other things being equal, what are the respective probabilities that each of the twelve pitch-classes is the novel chorale's key center?

Note that the model is not affected by the order in which it proceeds through the progressions or chorales in the training set. The context of each data point includes only the progression and the last note of its chorale melody. Recall, furthermore, that each progression has been encapsulated (assigned an arbitrary progression label). The system does not know what chords are in a progression, nor what its bass-line interval is – it only knows what the first bass note of the progression is, and what the progression's arbitrary label is.

A potential problem arises in the case of progression types appearing a very small number of times in the training set. Suppose a progression appears just once; when converted into a probability, this single observation becomes the hypothesis that the progression will appear on the same scale degree 100% of the time in novel chorales. Small numbers of observations lead to insupportably sharp probability distributions. To correct for this problem, a procedure called 'Laplace smoothing' is applied to the frequency tables before they are converted into probabilities; this simply involves adding 1 to the raw number of observations of each progression on each scale degree. When the actual number of observations is large, the smoothed probabilities are virtually identical to the unsmoothed probabilities. The smaller the number of actual observations, the more the smoothing step tends to redistribute probability mass evenly among all 12 scale degrees. In the single-observation case, the smoothed scale-degree probabilities are 15.4% (2/13) for the observed scale degree, and 7.7% (1/13) for each other scale degree.

3.2. Test phase

During the test phase, the system encounters each novel chorale in the form of an ordered list of chords, converting it into a list of encapsulated progression labels. The system proceeds through the list of chords once, maintaining a running estimate of global key center in the form of a probability-mass distribution over the twelve pitch classes.

The running estimate is taken into account as each new progression is read. The system is thus sensitive to the order in which it encounters progressions in the novel chorales of the test set, in contrast to the non-order-dependent character of the training phase. The algorithm is as follows:

1. *Initialization.* The system approaches each chorale with no assumption about key; it begins with a probability of 8.3 % (1/12) to all keys in its running estimate.
2. *Input.* The system reads a progression from the chorale, and retrieves the key-probability table for that progression that it generated during the training phase. Progressions not found in the training set are assigned equal-probability tables (8.3 % for all keys).
3. *Smoothing.* The previous running estimate is 'smoothed' by adding a constant amount of probability mass (5 %) to each key-center probability. This has the effect of decaying the system's memory, slightly biasing it toward revising its key-center estimate in the face of new data instead of holding on to previous estimates. A larger constant gives the system a more flexible sense of key, and a smaller constant makes it more conservative.
4. *Updating.* The system updates its running estimate of key probabilities by multiplying each probability in the smoothed running estimate (from Step Three) with the corresponding probability in the key-probability table for the current progression (from Step Two). The new running estimate is normalized, or scaled so that the total probability of 100 %.
5. *Looping.* If there are more progressions in the chorale, the system returns to Step Two. If not, the system moves on to the next chorale in the test set and returns to Step One.

4. Properties of the system's key-finding results

There were thirty-five chorales in the test set, and the system accurately identified the key center of all thirty-five after processing the last progression. The test set included three phrygian chorales, which Bach invariably ends with a half cadence; in each case, the system 'correctly' identified the last note of the melody (the phrygian final) as scale degree 5. This result shows a particular robustness since the system, by design, 'incorrectly' analyzed phrygian chorales in the training set – in the training phrase, scale degree 1 is defined as the last note in the chorale tune, contrary to Bach's phrygian-mode usage.

Let us now consider the system's behavior in some detail. Our discussion will focus on the system's analysis of the chorale "Befiehl du deine Wege," Riemenschneider no. 367 (see Example 2). As a phrygian chorale, this is one of the more difficult and problematic pieces in the test set, and it raises some intriguing questions about the broader implications for music theory of this highly specialized computational model.

4.1. Localness

More interesting than the system's final assessment of a chorale's key is what the running estimate looks like while the system is working its way through the chorale. The thick

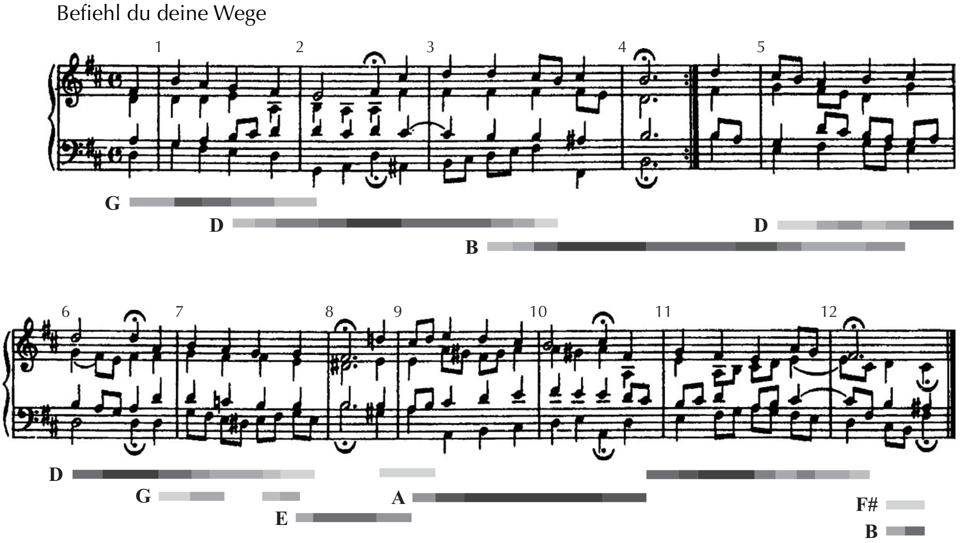

Example 2

shaded lines beneath the score in Example 2 indicate key centers with a probability of 20% or greater; darker parts of the line correspond to higher probabilities. In terms of what the system is literally trying to predict – the final note of the chorale tune – it fails utterly. Only the last two progressions bring about any serious prediction of F-sharp, and even then it is a distant second to B, the technically incorrect front-runner (remember that B is the phrygian final). But in terms of key-finding at the 'local' level, the system performs excellently, following along with all of Bach's modulations and tonicizations. The extent to which the system attends closely to local changes of key can be controlled at the smoothing step (step 3) of the algorithm. The more the system smoothes its running estimate (its 'memory'), the more susceptible it is to each incoming progression; a low or negative smoothing value will cause the system to hold more conservatively to its estimate over the long run. It is important to note that the system's memory is not at all particularized – it does not remember which progressions it has heard before, but only what its estimated key-center probabilities were. Put another way, it knows what it has thought, but it can't remember why.

4.2. Mehrdeutigkeit

The chorale begins with an ambiguous progression that can be read as either V–I in G or as I–IV in D. (In a different voice-leading environment, a passing seventh might distinguish these two cases, with C-sharp for D and C-natural for G.) The system, starting with a blank slate, initially estimates the key center as G and reads the bass line of the first progression as V–I. It admits D as a second possibility midway through the first bar, but only gives up on G at the arrival of the cadential progression II$^{6/5}$–V at the beginning of bar 2.

Upon taking the repeat (not shown in Example 1), the system holds on to the D-major interpretation of the first progression, reading the bass line of the first progression as I–IV.

What enables the system to analyze the same progression in different ways? Consulting Table 1, we see that all other things being equal, the progression in question (rank 3, n = 806), a descending-fifth pair of root-position major triads, is about 2.5 times more likely to be V–I than it is to be I–IV, yet these are both reasonably probable. The probability table is unable to distinguish between 'I–IV' and 'V–I of IV' (or, for that matter, between 'V–I' and 'I–IV of V'); and one would do well to wonder whether these do, in fact, mean different things. As Nicolas Meeùs writes, echoing Rameau: "Harmonic functions do not reside in chords, nor in the position of chords within an immanent tonal hierarchy. They result from a relation between chords. No chord is a dominant in itself, none is a tonic in itself; they become dominant and tonic with respect to each other when they occur in that relation" (2000, 15). For the system at hand, the only difference between what it hears at the beginning of each repeat of the 'Stollen' is its own predisposition: the first time through it hasn't got one, so it goes with the default, V–I [of IV of D], and the second time through it's predisposed to hearing D so it goes with the second choice, I–IV [of D itself].

4.3. Holism

The system holds on to its estimate of D despite what a human analyst would be likely to interpret as a move to B at the beginning of the second phrase. It is only a few chords into the phrase that the system even recognizes B as a possible key center, and only at the cadence formula that the system lets go of its D hypothesis. In the third phrase, the system's eagerness to return to D recalls its resistance to leave D in the second phrase. While this may seem to indicate a general bias toward the major key in a relative pair, there is scant evidence to support such a bias. In fact, for most standard dominant-tonic progression types, the system assigns an 'in-III' reading of the major-mode variant a higher probability than it assigns to the 'in-VI' reading of the minor-mode variant. This suggests that on the whole, the model trained on this particular training set actually his a minor-key bias.

Rather than indicating some abstract feature of the model, it is likely that the system's analysis of the second phrase is due instead to a much more specific bias: the three progressions that begin the second phrase, which result from the combination of a 9–8 suspension with a passing motion in the bass, are unusually strongly associated with the relative minor in the training set. In other words, this 'lick' is not simply a signifier of the minor mode in the Bach chorales, but more specifically a signifier of the relative minor of the major.[7] The model is sensitive to this kind of style- or corpus-specific detail. In fact, it is sensitive 'only' to this kind of detail: thanks to the encapsulation of chords and progres-

[7] Indeed, the model knows no distinction between major and minor at all. Its task is only to predict key center, operationalized as the last note of the melody. Modifying the model to predict mode would be possible, but given Bach's tendency to conclude minor-mode chorales with major triads, it would be a nontrivial task involving manual tagging of the training set or the limited use of pitch-class profiles during the training phase.

sions, any general 'biases' or 'tendencies' that we observe in its key-finding behavior can only be emergent properties of the statistical distribution of progressions in the corpus.

The holistic, emergent nature of the system is also evident when it comes to the most ambiguous portion of its analysis, the third and fourth phrases. For most of this passage the system admits two possibilities; it is only at the cadences that the analysis is clarified. Among all the progressions in the model, the most key-defining ones (those with the most sharply peaked transpositional distributions in the training set) are the progressions associated with cadences. In turn, they are key-defining in the model solely by virtue of appearing at a highly consistent transposition in the training set. Concepts of harmonic function, of dissonance treatment, of the imperfect seeking its perfection – to paraphrase Cohen (2003) paraphrasing Marchetto paraphrasing Aristotle – do not come into play here. All that matters is the bare fact that cadential progressions are less well-distributed transpositionally than are other kinds of progressions.

5. Questions

The surprising success of this model raises at least two families of questions. Particularly pressing are questions concerning the model's generalizability. The chorale repertoire is unique on a number of counts: it is relatively homorhythmic and texturally consistent; the harmonic rhythm is fast enough (relative to the frequency of note onsets) to guarantee a high information content in the harmonic domain; and it is easy to operationalize the identification of the crucial bass voice. Whether the system could be adapted to tonal repertoires that lack these characteristics – Classical piano sonatas, say, or Bach's solo violin suites – is unclear.

Another group of questions has to do with the challenges this computational model poses to the standard music-theoretic ways of characterizing the phenomenon of key. The key-finding system presented here, I have emphasized repeatedly, is designed to be radically naïve. It has a prodigious memory, like most statistical-learning models, and a peculiar talent for long-range hearing (insofar as it is able to detect intervallic relationships between a progression at the beginning of the chorale and the note at the end); but it has no capacity to engage in the kind of reductionist thinking favored by music theorists. Whatever knowledge the system has of pitch-class profiles, harmonic function, chord structure, dissonance, and so on is entirely implicit. The system has no access to explicit information about a chord progression other than its transpositional distribution in the training corpus, and yet it is able to reach an almost stunning degree of subtlety in its harmonic analysis of chorales it's never 'heard' before. This suggests that reductionist approaches to tonality may be off the mark, or at least that pitch-class reductionism might not be necessary for a principled account of key.

References

Chuan, Ching-Hua, and Elaine Chew. 2005. "Audio Key Finding Using FACEG: Fuzzy Analysis with the CEG Algorithm." *Music Information Retrieval Evaluation Exchange* (London, September 14).

Cohen, David. 2003. "'The Imperfect Seeks Its Perfection': Harmonic Progression, Directed Motion, and Aristotelian Physics." *Music Theory Spectrum* 23: 139–69.

Krumhansl, Carol. 1990. *Cognitive Foundations of Musical Pitch*. New York: Oxford University Press.

Meeùs, Nicolas. 2000. "Toward a Post-Schoenbergian Grammar of Tonal and Pre-tonal Harmonic Progressions." *Music Theory Online* 6.1.

Temperley, David. 1999. "What's Key for Key? The Krumhansl-Schmuckler Key-Finding Algorithm Reconsidered." *Music Perception* 17: 65–100.

——. 2002. "A Bayesian Approach to Key-Finding." In *Music and Artificial Intelligence*. Edited by C. Anagnostopoulou, M. Ferrand, and A. Smaill. Berlin: Springer: 195–206.

——. 2007. *Music and Probability*. Cambridge/MA: MIT Press.