Representing Harmony in Computational Music Cognition

Peter M. C. Harrison^{1, 2} & Marcus T. Pearce^{1, 3}

¹ Queen Mary University of London

² Max Planck for Empirical Aesthetics

³ Aarhus University

Author Note

Peter M. C. Harrison, School of Electronic Engineering and Computer Science, Queen Mary University of London; Marcus T. Pearce, School of Electronic Engineering and Computer Science, Queen Mary University of London.

This is an unpublished preprint that has yet to be peer reviewed (October 27, 2020). The manuscript is derived from part of Peter Harrison's doctoral thesis. The authors would like to thank Alan Marsden and David Baker for valuable feedback on an earlier version of the manuscript.

Peter Harrison was supported by a doctoral studentship from the EPSRC and AHRC Centre for Doctoral Training in Media and Arts Technology (EP/L01632X/1).

Correspondence concerning this article should be addressed to Peter M. C. Harrison, Max-Planck-Institut für empirische Ästhetik, Grüneburgweg 14, 60322 Frankfurt am Main, Germany. E-mail: peter.harrison@ae.mpg.de

Abstract

Cognitive theories of harmony require unambiguous formal models of how listeners internally represent chords and chord progressions. Previous modeling work often uses representation schemes heavily reliant on Western music theory, such as Roman-numeral and lead-sheet notation; however, we argue that such work should be complemented by models using representations that are closer to psychoacoustics and rely less on Western-specific assumptions. In support of this goal, we compile a network of 13 low-level harmonic representations relevant for cognitive modeling, organised into three symbolic, acoustic, and sensory categories. We implement this collection of representations in an easy-to-use object-oriented framework written for the programming language R and distributed in an open-source package called *hrep* (<http://hrep.pmcharrison.com>). We also discuss computational methods for deriving higher-level representations from these low-level representations. This work should ultimately help researchers to construct high-level models of harmony cognition that are nonetheless rooted in low-level auditory principles.

Keywords: auditory cognition; cognitive musicology; computational models; psychoacoustics

Representing Harmony in Computational Music Cognition

Cognitive science seeks to understand the mind in terms of mental representations and computational operations upon those representations. This computational metaphor is formalized by creating computational models of these cognitive representations and operations, which can then be tested against human behavior (Thagard, 2019).

Western music theory textbooks provide a natural starting point for designing computational cognitive representations of musical harmony. These textbooks teach chord categories such as V^7 , ii^7 , F^{sus4} , and D° , and describe how to derive such categories by inspecting a musical score or listening to a recording. These chord categories are intuitive to experienced music analysts, and they are used regularly by musicians who perform from lead sheets. In this context, it seems completely reasonable to have computational models operate over these Roman numeral or lead-sheet notations. Correspondingly, various researchers have created large corpora of harmonic transcriptions where the underlying representation is some variant of Roman numeral or lead-sheet notation (Broze & Shanahan, 2013; Burgoyne, 2011; Harte, Sandler, Abdallah, & Gómez, 2005; Neuwirth, Harasim, Moss, & Rohrmeier, 2018; Temperley & De Clercq, 2013), which have supported various valuable musicological and psychological studies of harmony (e.g., Temperley & De Clercq, 2013; Cheung et al., 2019; Mauch, Dixon, Harte, Casey, & Fields, 2007; Moss, Neuwirth, Harasim, & Rohrmeier, 2019).

However, when we began our own work studying various aspects of harmony cognition, we felt uneasy beginning with these kinds of representations. First, we wanted psychoacoustic considerations (e.g., roughness, harmonicity, spectral similarity) to play a deep role in our modeling, yet these kinds of psychoacoustic features are very distant from Roman-numeral or lead-sheet representations. Second, we were wary of how these notation systems embody implicit assumptions from Western common-practice music; without disputing the value of these notation systems in this domain, we rather wished to develop

models that made fewer style-specific assumptions and could in theory generalize cross-culturally. Third, we felt that, to the extent that our ultimate goal was a comprehensive computational model of harmony cognition, this model should begin earlier in the perceptual process: it should not take music-theoretic chord categories for granted, but should rather explain how these chord categories are derived in the mind of the listener.

Various parts of the literature suggested lower-level harmonic representations that seemed more appropriate as a starting point. Pitch-class sets, commonly used for atonal music analysis, provided one possible way to avoid assumptions about Western tonality. A second possibility was the OPTIC framework of Callender, Quinn, and Tymoczko (2008), which defines a set of 32 representations embodying different combinations of cognitive equivalence principles such as octave equivalence and transposition equivalence. Likewise, the idealized harmonic spectra used by Plomp and Levelt (1965) seemed useful for modeling psychoacoustic interactions between partials, and the smoothed harmonic spectral of Milne, Sethares, Laney, and Sharp (2011) seemed useful for modeling psychoacoustic relationships between chords. Other important models of tonal cognition (Collins, Tillmann, Barrett, Delbé, & Janata, 2014; Leman, 2000) instead required raw audio as the input. It became apparent that it would not be sufficient to adopt a single such representation for all modeling work; instead, one would need to move flexibly between different representations according to the different models being applied.

This paper describes our efforts to facilitate working with these different representations. We review what we term ‘low-level’ harmonic representations that are closely linked to the psychoacoustic side of music perception, and precede the more subjective and style-dependent aspects of cognitive representation such as key finding, root finding, and functional classification. We organize this review in terms of a small collection of basic perceptual principles, such as octave invariance, transposition invariance, and perceptual smoothing, and we explain different use cases for the constituent

representations, as well as showing how they can be derived from existing music corpora. We present this review alongside an open-source R package, *hrep*, that implements these different representations and encodings in a generalizable object-oriented framework.¹ Our aim is that this framework should ultimately help researchers to construct high-level models of harmony cognition that are nonetheless rooted in low-level auditory principles.

Related work

Several large corpora of harmonic transcriptions are available in the literature, including the iRb corpus (Broze & Shanahan, 2013), the Billboard corpus (Burgoyne, 2011), the Annotated Beethoven Corpus (Neuwirth et al., 2018), the rock corpus of Temperley and De Clercq (2013), and the Beatles corpus of Harte et al. (2005). The iRb corpus (Broze & Shanahan, 2013) and the Billboard corpus (Burgoyne, 2011) express chords using letter names for chord roots and textual symbols for chord qualities; for example, a chord might be written as ‘D:min7’ where ‘D’ denotes the chord root and ‘min7’ denotes a minor-seventh chord quality.² The Beatles corpus of Harte et al. (2005) uses a similar representation, except with chord qualities expressed as explicit collections of pitch classes expressed relative to the chord root, with for example $D\#:(b3, 5, b7)$ denoting a $D\#$ minor-seventh chord. The rock corpus of Temperley and De Clercq (2013) and the Annotated Beethoven Corpus of Neuwirth et al. (2018) provide deeper levels of analysis: they express chords relative to the prevailing tonality, and provide functional interpretations of the resulting chord progressions, writing for example ‘V/vi’ to denote the secondary dominant of the submediant.

These preceding representations rely on Western-specific notions of chord roots and chord qualities. The General Chord Type representation of Cambouropoulos (2016) is

¹ <http://hrep.pmcharrison.com>; <https://doi.org/10.5281/zenodo.2545770>.

² A *chord quality* defines a chord’s pitch-class content relative to the chord root. See the *Pitch-class set* section for a definition of ‘pitch class’ and the *Chord roots* section for a definition of ‘chord root’.

designed to generalize better to non-Western musical styles, while still keeping some of the useful analytic qualities of the original representations. The representation is parametrized by a style-dependent ‘consonance vector’ that quantifies the relative consonance of different intervals within the musical style. A useful advantage of the representation scheme is that it is fully algorithmic. For each chord, the algorithm takes as input a collection of pitch classes, and organizes these pitch classes into three components: one pitch class labeled as the root, a set of pitch classes labeled as the chord’s base, and the leftover pitch classes which are labeled as the chord’s extensions. Applied to non-Western tonal systems, the algorithm produces harmonic representations that seem to work well for statistical modeling and music generation (Cambouropoulos, Kaliakatsos-Papakostas, & Tsougras, 2014). However, the algorithm is only loosely motivated by human cognition, and has yet to receive systematic perceptual validation.

The OPTIC framework of Callender et al. (2008) seems to achieve a good combination of strong cognitive foundations and cross-style generalizability. This framework defines 32 different harmonic representations in terms of different combinations of invariances, namely (a) octave invariance, (b) transposition invariance, (c) permutation invariance, (d) inversion (or reflection) invariance, and (e) cardinality invariance. All (except perhaps inversion invariance) can be strongly motivated by perceptual principles, and do not rely substantially on Western music theory (though see Jacoby et al., 2019 for evidence that octave invariance is not a cross-cultural universal). Moreover, the framework encompasses several traditional music-theoretic representations such as pitch classes, pitch class sets, and chord types. However, it is insufficient for our own work because it does not offer the kinds of acoustic and sensory representations required for psychoacoustic modeling (e.g., Hutchinson & Knopoff, 1978; Milne & Holland, 2016).

Other representations have been developed to characterize the intervallic relationships between successive chords, such as the voice leadings described by Tymoczko (2008), the

voice-leading types of Quinn and Mavromatis (2011) and Sears, Arzt, Frostel, Sonnleitner, and Widmer (2017), the interval function of Lewin (1987), and the directed interval class vector of Cambouropoulos (2016). Such representations are indisputably valuable for music analysis, but they fall more in the realm of voice leading (connecting the notes of successive musical chords to form melodies) than of harmony, the topic of the present work. We will return to this topic in the *Deriving higher-level representations* section.

Proposed representations

Here we present our proposed collection of low-level harmony representations. These representations are organized into three categories: the *acoustic*, the *sensory*, and the *symbolic*.³ Acoustic representations characterize musical sound in a listener-free manner. Sensory representations reflect listeners’ perceptual images of these sounds. Symbolic representations are succinct and human-readable descriptions of musical chords that ideally have a strong cognitive basis.

When describing such a framework, it is important to be clear concerning our definitions of terms such as ‘harmony’ and ‘chord’. Broadly speaking, harmony may be defined as the vertical aspect of music, concerning the way in which notes played in close temporal proximity can be integrated into meaningful units termed ‘chords’, and the way that these chords are organized sequentially to form chord progressions. However, the definition of what constitutes a chord requires some nuance. In many musical styles, the chord is not written out transparently in the score, but rather has to be inferred from the musical surface, a subjective process that often requires taking account of musical context. A chord is therefore better defined as a cognitive category rather than an acoustic object. This cognitive process of extracting chords from the musical surface may be intuitive for the music theorist, but it is not yet well-understood from a cognitive perspective (though

³ These categories were inspired by Babbitt’s (1965) graphemic, acoustic, and auditory representational categories.

see Rohrmeier & Cross, 2008; Barthélemy & Bonardi, 2001; Chen & Su, 2018; Chuan & Chew, 2011; Ju, Condit-Schultz, Arthur, & Fujinaga, 2017; Kröger, Passos, Sampaio, & Cidra, 2008; Masada & Bunesu, 2017; Mearns, 2013; Pardo & Birmingham, 2002; Raphael & Stoddard, 2004; Temperley & Sleator, 1999; and Winograd, 1968 for relevant computational models). In the present work we avoid this unsolved problem of cognitive extraction; instead, we make the simplifying assumption that we are dealing with unelaborated block chords.

Following Callender et al. (2008), our representations are organized around the notion of invariance. Invariance means that a chord retains some kind of musical identity under a given operation. For example, transposition invariance means that a chord retains its identity when all its pitches are transposed by the same pitch interval; octave invariance means that a chord retains its identity when individual pitches are transposed by octave intervals.

An invariance principle may also be formulated as an *equivalence relation*: saying that a chord is invariant under a given operation is the same as saying that a chord is equivalent to all other chords produced by applying the operation. Equivalence relations partition sets into *equivalence classes*, sets of objects that are all equal under some equivalence relation. If we label each object by its equivalence class, we produce a new representation that embodies the original invariance principle behind the equivalence relation. For example, if we begin with the seven triads of the C major scale (C major, D minor, E minor, ...) and add transposition invariance, we get three equivalence classes: one of major triads (C major, F major, G major), one of minor triads (D minor, E minor, A minor), and one of diminished triads (B diminished). These equivalence classes therefore define a new representation with three transposition-invariant chord qualities: major, minor, and diminished. See Callender et al. (2008), Tymoczko (2011), and Lewin (1987) for further discussions of equivalence classes induced by musical invariances.

The invariance principles described here each have a well-established perceptual basis. Octave invariance describes how listeners perceive tones separated by octaves to share some essential quality termed *chroma* (Bachem, 1950). Transposition invariance reflects the phenomenon of relative pitch perception (Plantinga & Trainor, 2005). Tones with different spectra but identical fundamental frequencies share an invariant perceptual quality termed *pitch* (Stainsby & Cross, 2009). Note that these invariance principles are not necessarily equally strong in all musical styles or cultural groups (Jacoby et al., 2019), but they are certainly well-established for Western listeners.

We organize our different low-level representations as a network (Figure 1). This network clarifies the sequence of computational operations required for translating between different representations, and shows how these representations may be organized according to different perceptual invariances. This structure is made more explicit in Table 1, which lists the invariance principles that apply to each representation, their alphabet sizes assuming a 12-tone scale, and the corresponding class names in the *hrep* package.

Figure 1. A network of low-level harmony representations, organized into three representational categories: symbolic, acoustic, and sensory. Arrows between representations indicate well-defined computational translations between representations.

Table 1
The 13 low-level harmony representations with their alphabet sizes and invariances.

Name		Invariance				
Textual	Computer	Alphabet size	Tone spectra	Tuning	Inversion	Octave Transposition
Symbolic						
Pitch chord	pi_chord		✓	✓		
Pitch-class set	pc_set	4,095	✓	✓	✓	✓
Pitch-class chord	pc_chord	24,576	✓	✓		(✓)
Pitch chord type	pi_chord_type		✓	✓		✓
Pitch-class chord type	pc_chord_type	2,048	✓	✓		(✓)
Pitch-class set type	pc_set_type	351	✓	✓	✓	✓
Acoustic						
Frequency chord	fr_chord		✓			
Sparse pitch spectrum	sparse_pi_spectrum					
Sparse pitch-class spectrum	sparse_pc_spectrum					✓
Sparse frequency spectrum	sparse_fr_spectrum					
Waveform	wave					
Sensory						
Smooth pitch spectrum	smooth_pi_spectrum					
Smooth pitch-class spectrum	smooth_pc_spectrum					✓

Note. Here ‘inversion’ refers to changing the chord’s bass pitch class while maintaining its pitch-class set. ‘Alphabet size’ refers to the total number of possible categories in the representation, assuming chord pitches drawn from a 12-tone scale.

Symbolic representations

It is hard to find a good yet comprehensive naming system for symbolic chord representations. Standard terminology in the literature (e.g., ‘pitch-class set’, ‘chord type’) fails to cover all of our required representations. Callender et al. (2008) provide an elegant name-generation system based on invariance properties, but the resulting names are unintuitive unless one has memorized the system. Our own naming system is instead defined by the following conventions:

1. ‘Pitch-class’ representations express octave invariance;
2. ‘Chord’ representations identify which pitch class corresponds to the lowest pitch in the chord, termed the *bass note*;
3. ‘Type’ representations express transposition invariance.

The six symbolic representations defined here are the *pitch chord*, the *pitch-class set*, the *pitch-class chord*, the *pitch chord type*, the *pitch-class chord type*, and the *pitch-class set type*. Many of these representations have a basis in prior literature. Table 2 displays these six representations as derived for several example chords. We will now introduce each representation in detail.

Pitch chord. The *pitch chord* representation is the most granular of the symbolic representations considered here. The other five symbolic representations are defined by partitioning this representation into equivalence classes. The representation might equally be termed ‘pitch set’, but we use the term ‘pitch chord’ for consistency with the conventions described above. Each pitch chord is expressed as a finite and non-empty set of pitches; each pitch is represented as a MIDI note number, where 60 corresponds to middle C (C4, c. 262 Hz). Integers correspond to pitches in 12-tone equal temperament; however, non-integer pitches (and later, pitch classes) are also permitted so as to express different

Table 2

Symbolic representations for three example chords: (A) C major triad, 1st inversion; (B) C major triad, second inversion; (C) E dominant seventh, third inversion.

Representation	Chord		
	A	B	C
pi_chord	{52, 60, 67}	{43, 52, 60}	{50, 59, 64, 68}
pc_set	{0, 4, 7}	{0, 4, 7}	{2, 4, 8, 11}
pc_chord	(4, {0, 7})	(7, {0, 4})	(2, {4, 8, 11})
pi_chord_type	{0, 8, 15}	{0, 9, 17}	{0, 9, 14, 18}
pc_chord_type	{0, 3, 8}	{0, 5, 9}	{0, 2, 6, 9}
pc_set_type	{0, 4, 7}	{0, 4, 7}	{0, 3, 6, 8}

scale and tuning systems.⁴ A chord containing n unique notes is then written as a set of n numbers; for example, the set {54, 62, 69} represents a D major triad in first inversion comprising the notes F#3, D4, and A4. Transposition then corresponds to simple addition. Note that duplicate pitches are not encoded, reflecting how such pitches tend to be perceptually fused in the mind of the listener.

Pitch chord type. *Pitch chord types* express transposition invariance, and are defined by the following equivalence relation: chords are equivalent if they can be transposed to the same pitch chord representation. Pitch chord types are represented as sets of numbers, where each chord note is represented as its pitch interval from the bass note. Given a pitch chord, the pitch chord type may be computed by simply subtracting the bass note from all pitches: for example, the pitch chord type of the pitch chord {54, 62, 69} is {0, 8, 15}. This representation is used by certain transposition-invariant consonance models, for example Gill and Purves (2009) and Stolzenburg (2015).

Pitch-class set. *Pitch-class sets* express the principle of octave invariance. The representation is defined by partitioning the set of pitch chords under the following

⁴ Though some theorists might conceptualize pitches and pitch classes as fundamentally categorical entities, here we follow Tymoczko (2011) and Milne et al. (2011) in treating pitches and pitch classes as continuous objects, which may subsequently be classified by the listener into discrete categories.

equivalence relation: two chords are equivalent if they comprise the same pitch classes. Pitch classes are meanwhile defined as equivalence classes over pitches: two pitches correspond to the same pitch class if they are separated by an integer number of octaves. A pitch class may be represented as a number between 0 (inclusive) and 12 (exclusive), computed by applying the modulo 12 operator to MIDI note numbers; a pitch-class set is then a collection of pitch classes (discarding duplicates). Pitch classes (and hence pitch-class sets) can be transposed by adding x to each number modulo 12. The pitch-class set representation is commonly used for atonal music analysis (e.g., Rahn, 1980), and sometimes also for tonal music analysis (e.g., Rohrmeier & Cross, 2008; Harrison & Pearce, 2018a). An example pitch-class set would be the A minor seventh chord $\{0, 4, 7, 9\}$.

Pitch-class set type. The *pitch-class set type* representation combines the invariances of the pitch chord type representation and the pitch-class set representation, namely transposition invariance and octave invariance. The representation may be formally defined by the following equivalence relation: chords are equivalent if they can be transposed to the same pitch-class set representation. The computational challenge in defining the pitch-class set type representation is to find a deterministic algorithm that returns a single representative of the equivalence class corresponding to the above equivalence relation given any member of that equivalence class. This can be achieved through two steps: (a) find the normal order of the chord’s pitch-class set using Rahn’s (1980) algorithm,⁵ and (b) transpose this normal-order pitch-class set such that the first element is 0. For example, the pitch-class set $\{0, 7, 9\}$ has a normal order of $(7, 9, 0)$, and so its pitch-class set type is $\{0, 2, 5\}$. Examples of pitch-class set types include the major triad $\{0, 4, 7\}$ and the augmented triad $\{0, 4, 8\}$.

Pitch-class chord. *Pitch-class chords* express a limited form of octave invariance that differentiates chords with different bass pitch classes. Musicians use the term *inversion* to describe the process of changing a chord’s bass pitch class while keeping the pitch-class

⁵ See Appendix A for a definition of Rahn’s (1980) algorithm.

set the same; we might therefore say that pitch-class chords lack inversion invariance. The identity of the bass pitch class is considered important in music theory, and is encoded in both figured-bass and Roman numeral notation schemes. Correspondingly, we might hypothesize that the identity of the bass pitch class contributes significantly to a chord’s perceptual identity, though admittedly there is a surprising lack of cognitive studies testing this empirically. Nonetheless, this special role of the bass note is already acknowledged by the music information retrieval literature, for example by Mauch and Dixon (2008) who introduced dedicated bass chromagrams into acoustic chord modeling. The pitch-class chord representation may be formally defined by the following equivalence relation: two chords are equivalent if a) they have the same pitch-class set and b) they have the same bass pitch class, with the bass pitch class being defined as the pitch class of the lowest chord pitch. Correspondingly, a pitch-class chord may be represented as a tuple of its bass pitch class and the set of non-bass pitch classes; for example, the pitch chord $\{54, 62, 69, 74\}$ can be represented as the pitch-class chord $(6, \{2, 9\})$, and correspondingly labeled as a D major triad in first inversion. Like pitch-class sets, pitch-class chords may be transposed by x semitones by adding x to each number modulo 12.

Pitch-class chord type. The *pitch-class chord type* representation combines the invariances of the pitch chord type representation and the pitch-class chord representation: it expresses both transposition invariance and a limited form of octave invariance that differentiates chords with different bass pitch classes. The representation may also be defined by the following equivalence relation: chords are equivalent if they can be transposed to the same pitch-class chord representation. Pitch-class chord types are represented as sets of numbers, where each chord note is represented as its pitch-class interval from the bass pitch class.⁶ An example of a pitch-class chord type would be the first-inversion major triad $\{0, 3, 8\}$.

⁶ The pitch-class interval from x to y is defined as $y - x \bmod 12$.

Integer encodings. When analysing the statistics of chord progressions, it is often useful to define a system for encoding each chord as a single integer. For example, these integers could define the indices of a Markov transition matrix, or the coding of a one-hot input vector for a neural network. A classic example of a harmonic alphabet is Roman numeral notation, which encodes each chord as the scale degree of its root. Another example of a musical alphabet is the ‘Forte number’, which indexes the prime forms of pitch-class sets.⁷ Appendix B provides simple encoding algorithms for four of our symbolic representations: pitch-class sets, pitch-class set types, pitch-class chord types, and pitch-class chords. Note that these algorithms each assume that the input chord is quantized to the 12-tone scale, although they could be straightforwardly generalized to arbitrary subdivisions of the octave.

Potential extensions. Each of the symbolic representations conveys a binary notion of presence: a chord component (e.g., a pitch or pitch class) may be absent or present, but nothing in between. It could be useful to incorporate a more graduated notion of presence, to represent the fact that given chord components might be weighted more than others, perhaps by doubling pitch classes, playing notes more loudly, or situating notes in more perceptually salient registers. It could also be useful to differentiate newly sounded notes from notes that are sustained from a previous chord. Both of these types of metadata could be incorporated as numeric or Boolean vectors with the same length as the original note vector. However, we do not pursue these possibilities further in the present work.

Acoustic representations

Here we describe five acoustic representations: the *frequency chord*, the *sparse frequency spectrum*, the *sparse pitch spectrum*, the *sparse pitch-class spectrum*, and the waveform.

⁷ See Carter (2002) and Martino (1961) for related efforts.

Frequency chord. Each *frequency chord* is represented as a set of positive real numbers, corresponding to the fundamental frequencies of the chord’s constituent tones. These frequencies can be approximated by the formula

$$f = f_{ref} \times 2^{(p-69)/12}, \quad (1)$$

where f is the frequency to be computed in Hz, $f = f_{ref}$ is the reference frequency (typically 440 Hz), and p is the input pitch, expressed as a MIDI note number. This formula can be also adjusted to incorporate stretched octaves, for example by applying $f = f_{ref} \times 2^{(p-69)/\gamma}$ (Parncutt & Strasburger, 1994) or $f = f_{ref} \times \gamma^{(p-69)/12}$ (Sethares, 2005) with γ taking a value slightly below 12 or slightly above 2 respectively.

Sparse frequency spectrum. The *sparse frequency spectrum* is generated from the frequency chord by incorporating knowledge of the spectral content of the tones used to play the chord (Figure 2A). Each chord tone is modeled as a *complex tone*, defined as a superposition of pure tones, each of known frequency and amplitude. This definition ignores phase as well as any temporal evolution of the chord tones. In particular, we default to modeling each tone as a *harmonic complex tone*, where the i th partial has a frequency i times that of the fundamental frequency and an amplitude that decreases steadily with increasing i . The i th partial is then termed the i th harmonic. Similar to previous literature (e.g., Hutchinson & Knopoff, 1978; Milne & Holland, 2016; Parncutt & Strasburger, 1994), we default to modeling each complex tone with 11 partials, and give the i th partial an amplitude of $1/i$, with the amplitudes being provided in arbitrary units. These defaults are somewhat arbitrary, and the researcher is encouraged to adjust the number of harmonics and their amplitudes to reflect the particular sound being modeled.

Once each chord tone is decomposed into its partials, these partials must be combined together into one spectrum. For this combination process, it is common to define

Figure 2. Acoustic and sensory representations for a C major chord in first inversion comprising the MIDI pitches $\{52, 60, 67\}$, synthesized with 11 harmonics with amplitude roll-off $1/n$. (A) Sparse frequency spectrum; (B) sparse pitch spectrum; (C) sparse pitch-class spectrum; (D) waveform; (E) smooth pitch spectrum; (F) smooth pitch-class spectrum.

a procedure for combining partials from different tones with (approximately) the same frequency. Such a procedure requires two decisions: (a) how close do two partials need to be before they will be combined, and (b) what is the amplitude of the resulting partial. The latter question is easy to answer: in most musical instruments, different chord tones will be produced by different oscillators, so they will have incoherent phases,⁸ and it is a standard result that the incoherent superposition of two waves with (approximately) identical frequencies and amplitudes x, y yields a resultant amplitude of

⁸ Incoherence means that the phases of the two waves have no systematic relationship.

$$(x^2 + y^2)^{0.5}. \quad (2)$$

The former question has a less consistent answer in the literature. Here we adopt the heuristic of expressing each frequency as a (possibly non-integer) MIDI note number, rounding the MIDI note number to six decimal places, combining any partials with identical results, and translating the MIDI note numbers back to frequencies. Six decimal places should be sufficiently precise to differentiate non-identical partials in most practical contexts, but sufficiently lax to be robust to floating-point inaccuracies in software implementations. Alternatively, one could adopt a perceptually motivated criterion related to perceptual discrimination thresholds.⁹

Ideally, computational models taking sparse frequency spectra as input should be invariant, in the limit, to the arbitrary decision of whether or not to combine two partials with almost-identical frequencies; unfortunately, this is seldom true in practice. For example, Dillon (2013) explains how several historic consonance models (e.g., Hutchinson & Knopoff, 1978; Sethares, 1993, 2005) fail to express invariance to partial combination, and presents an alternative consonance model that captures this desired invariance. However, this model has yet to be systematically tested, and it is therefore unclear how much the invariance principle matters in practice.

Sparse pitch spectrum. The *sparse pitch spectrum* corresponds to a version of the sparse frequency spectrum where partial frequencies are converted to MIDI note numbers (Figure 2B). This logarithmic transformation provides a more intuitive way of visualising how spectral content maps to 12-tone equal temperament, and provides a better account of the subjective notion of pitch distance (Jacoby et al., 2019). The mapping between frequency and MIDI note number is defined by Equation 1.

⁹ For example, Moore, Peters, and Glasberg (1985) report a frequency discrimination threshold of approximately 1% for the partials in a harmonic complex tone.

Sparse pitch-class spectrum. The *sparse pitch-class spectrum* corresponds to an octave-invariant version of the sparse pitch spectrum (Figure 2C). It can be computed from the sparse pitch spectrum by expressing each pitch as a pitch class (using the modulo 12 operation) and combining partials with identical pitch classes. As before, pitch classes are rounded to six decimal places before testing for equality, and amplitudes are combined assuming incoherent summation (Equation 2).

It is also possible to compute pitch-class spectra from pitch-class sets, because both representations are invariant to octave transpositions of chord tones. This may be achieved by constructing a pitch chord containing one instantiation of every pitch class in the pitch-class set, and then computing the sparse pitch spectrum and hence the sparse pitch-class spectrum as described above.

The sparse pitch-class spectrum is closely related to the *chroma vector* representation commonly used in music audio analysis (e.g., Collins et al., 2014; Lartillot, Toivainen, & Eerola, 2008); both are octave-invariant representations that capture the perceptual notion of chroma. The primary difference between the two is that the former represents the pitch-class domain as a continuous space, whereas the latter partitions the pitch-class domain into 12 categories. The former approach is preferable for an acoustic representation because it avoids imposing the cultural assumption of a specific scale system.

Waveform. The *waveform* characterizes the chord as the fluctuation of sound pressure over time (Figure 2D). It may be computed from the sparse frequency spectrum using additive synthesis, where separate sine waves are computed for each partial in the spectrum, and then additively combined. This operation may be approximately reversed by conducting a Fourier transform and then applying a spectral peak-picking algorithm (see e.g., Essentia toolbox, Bogdanov et al., 2013; Mirtoolbox, Lartillot et al., 2008).

Sensory representations

Smooth pitch/pitch-class spectrum. In order to simulate how perceptual inaccuracies reduce the effective resolution of pitch perception, computational implementations of sensory representations can apply pitch-domain *smoothing*. Milne and colleagues have shown that a simple model of this smoothing effect is remarkably effective for reproducing a variety of cognitive phenomena (e.g., diatonic tonality, Milne, Laney, & Sharp, 2015; harmonic distance, Milne & Holland, 2016; perceived change, Dean, Milne, & Bailes, 2019). Our smooth spectra representations are computed from sparse spectra using this smoothing technique. The sparse spectrum is first expressed as a *dense spectrum*, a numeric vector where each bin corresponds to a pitch window of one cent (i.e., a hundredth of a semitone) and each bin value corresponds to the spectral amplitude in that pitch window. This is achieved by initialising a 0-indexed vector of zeros with 1200 elements per octave, where the 0th element corresponds to 0 cents, the 1st element corresponds to 1 cent, and so on, and then iterating through every partial in the sparse spectrum, rounding its frequency to the nearest cent, and incrementing the corresponding element in the vector using the incoherent amplitude summation rule (Equation 2). To produce the smooth spectrum, we then convolve this dense spectrum with a Gaussian function with unit mass and standard deviation of 10 cents, after Dean et al. (2019). In the case of pitch-class spectra, this should be a circular convolution, so that the smoothing wraps round the extremes of the pitch-class vector. Perceptual similarity between smooth spectra can then be simulated using geometric similarity measures such as cosine similarity:

$$s(\mathbf{x}, \mathbf{y}) = \frac{\sum_i x_i y_i}{\sqrt{\sum_j x_j^2 \sum_k y_k^2}} \quad (3)$$

where $s(\mathbf{x}, \mathbf{y})$ is the cosine similarity between vectors $\mathbf{x}$ and $\mathbf{y}$, and x_i, y_i denote the i th elements of $\mathbf{x}$ and $\mathbf{y}$ respectively. This method can be used to produce two types of smooth

spectrum: the *smooth pitch spectrum* (Figure 2E), and the *smooth pitch-class spectrum* (Figure 2F). Both pitch (Dean et al., 2019; Milne et al., 2016) and pitch-class versions (Milne & Holland, 2016; Milne et al., 2015) have proved useful in prior cognitive research.

Several subtleties should be acknowledged when implementing these spectral smoothing techniques. First, the original presentation of the technique treated the smoothed spectrum as an ‘expectation vector’, with each element being interpreted as the expected number of tones that the listener should hear at that pitch or pitch class (Milne et al., 2011). However, this probabilistic interpretation has limited empirical basis, and subsequent work often prefers a less specific ‘perceptual weight’ interpretation. Second, Milne assumed that the i th harmonic of a harmonic complex tone should additively contribute a value of $1/i^\rho$ to this perceptual weight spectrum, where ρ is a numeric parameter optimized to the data. This assumption of additive combination in perceptual space is difficult to reconcile with the formula for the acoustic combination of incoherent partials (Equation 2). However, in later work Milne and colleagues removed this inconsistency by applying the smoothing directly to the acoustic amplitude spectrum (Dean et al., 2019). We prefer this latter approach for its applicability to arbitrary acoustic spectra, but implement both approaches in the *hrep* package.

As is apparent from Figure 1, the computational translation from pitch-class sets to smooth pitch-class spectra is not well-defined. This is because pitch-class sets discard information about the number of chord pitches representing each pitch class, information which is retained implicitly in the smooth pitch-class spectrum. However, if necessary a representative spectrum can still be computed by assuming that each pitch class is assigned exactly one chord pitch.

Further alternatives and omissions

We have purposefully omitted the ‘prime form’ representation from pitch-class set theory, which expresses octave invariance, transposition invariance, and reflection invariance (Forte, 1973; Rahn, 1980).¹⁰ Our reasoning is that there is little evidence for reflection invariance in harmony perception: for example, the major and minor triads are reflections of each other, yet they imply opposite emotional valences to Western listeners (‘happy’ versus ‘sad’ respectively).¹¹

We have also purposefully omitted several representations commonly used to study voice leading. A *multiset* generalises the notion of pitch chords and pitch-class sets to differentiate between chords containing different numbers of the same element, for example $\{C, C, E, G\}$ (e.g., Tymoczko, 2008). A *voice leading* describes how each element of one chord moves to an element in the next chord, for example $\{F \rightarrow E, A \rightarrow G, C \rightarrow C\}$ (Tymoczko, 2008), and a *voice-leading type* is a transpositionally equivalent set of voice leadings (Quinn & Mavromatis, 2011; see Sears et al., 2017 for a related definition).¹² By omitting these voice-leading representations, we acknowledge the conventional distinction between harmony (the construction and arrangement of chords) and voice leading (connecting the constituent notes of successive musical chords to form simultaneous melodies).

Our sensory representations are based on the smoothed pitch- and pitch-class spectra of Milne et al. (2011). However, several other researchers have formulated sensory

¹⁰ Pitch-class set theorists typically use the word ‘inversion’ instead of ‘reflection’, but we use the latter to avoid confusion with the unrelated concept of inversion in tonal harmony. For example, suppose that we wish to reflect the major-triad pitch-class set $\{0, 4, 7\}$ about the pitch-class 4: we have $0 \rightarrow 8$, $4 \rightarrow 4$, and $7 \rightarrow 1$, producing the minor-triad pitch-class set $\{1, 4, 8\}$.

¹¹ However, see Dienes and Longuet-Higgins (2004) and Krumhansl, Sandell, and Sergeant (1987) for studies of inversion salience in the perception of serially presented tone rows.

¹² Sears et al.’s (2017) representation only captures the voice leading between successive bass notes, and expresses each successive chord as a set of pitch-class intervals above the bass note; it may be considered a hybrid between harmonic and voice-leading approaches.

representation schemes that can be applied to harmony modeling. Parncutt and Strasburger (1994) present a psychoacoustic model of music perception, based on Terhardt’s (1982) pitch perception model, that incorporates features such as auditory masking and spectral-domain pitch detection. We have implemented this model in the freely available R package *parn94*.¹³ This model can be used as an alternative to the smooth pitch spectrum representation described here. Leman (2000) provides an alternative psychoacoustic model of pitch perception, which combines Van Immerseel and Martens’ (1992) simulation of the auditory periphery with a time-domain pitch detection algorithm. The model takes an audio signal as input, and returns a time-varying vector of activations for different pitch windows; this representation can then be used as an alternative for the smooth pitch spectrum representation.¹⁴ Collins et al. (2014) implement two further representations that derive from Leman’s (2000) representation: a chroma vector representation and a tonal space representation. The chroma vector representation adds octave invariance, whereas the tonal space representation projects Leman’s representation onto a self-organizing map in the shape of a torus, intended to capture the topological structure of Western tonality.¹⁵ Each of these models (Collins et al., 2014; Leman, 2000; Parncutt, 1989) add significant complexity to the process of deriving sensory representations; future work needs to examine the extent to which this added complexity is useful for characterising harmony cognition.

Deriving higher-level representations

By focusing so far on low-level representations, we do not mean to diminish the importance of higher-level representations such as chord roots and functional chord

¹³ <https://github.com/pmcharrison/parn94>; <https://doi.org/10.5281/zenodo.2545759>.

¹⁴ At the time of writing, Leman’s (2000) model was available in the IPEM toolbox, hosted at <https://github.com/IPEM/IPEMToolbox>.

¹⁵ At the time of writing, the implementation of Collins et al. (2014) was available in the Janata Lab Music Toolbox, hosted at <http://atonal.ucdavis.edu/resources/software/jlmt/>.

categories for cognitive modeling. Rather, we intend the low-level representations to provide the initial step in computational cognitive models simulating the extraction of such higher-level representations. To support future readers in constructing such pipelines, we now provide a brief bibliographic survey of current computational techniques for deriving higher-level harmonic representations.

Voice leading

Chords played in succession can imply latent melodic lines, even if these melodic lines are not made explicit by performance characteristics such as timbre and loudness. Under Bregman’s (1990) terminology of auditory scene analysis, we can say that successive notes in chord progressions are grouped into parallel auditory streams, where individual streams cohere on the basis of pitch proximity between adjacent notes. This organization into auditory streams, or melodic lines, may be termed *voice leading*.

We may wish to simulate this process of inferring voice leading from a chord sequence. Tymoczko (2006) has presented an efficient algorithm for computing voice leadings within pairs of chords; this algorithm deterministically finds a voice leading that minimizes the aggregate distance moved by the individual melodic lines. In its original presentation, the algorithm finds minimal voice leadings between pitch-class sets, but it can trivially be modified to find minimal voice leadings between pitch chords. We have implemented this algorithm in the freely available R package *minVL*.¹⁶ Building on this work, we have also developed an algorithm for generating voice leadings for longer chord sequences while balancing multiple constraints such as maximizing voice-leading efficiency, minimizing interference between partials, and avoiding parallel octaves and fifths (Harrison & Pearce, 2020b).

¹⁶ <https://github.com/pmcharrison/minVL>

Chord roots

Chord roots have been considered central to harmonic analysis since Rameau (1722). A chord root may be defined as a single pitch class that summarizes the tonal content of a given chord. Chords are often notated with reference to these roots: for example, lead sheets typically represent chords by combining the root pitch class (written in letter-name notation, e.g., ‘D’) with a textual label expressing the other pitch classes relative to the root (e.g., ‘maj7’ implies that the chord contains pitch classes 0, 4, 7, and 11 semitones above the root). Roman numeral notation also uses chord roots, which are typically expressed relative to the local tonic (see the section *Tonality*).

It has been proposed that chord roots have a psychological basis in pitch perception (Parncutt, 1988; Terhardt, 1974, 1982). Terhardt and Parncutt argue that listeners infer fundamental frequencies of acoustic spectra by comparing these spectra to harmonic templates with different fundamental frequencies. Candidate fundamental frequencies are perceived as *virtual pitches*, which Terhardt and Parncutt equate with chord roots. Recent empirical work indeed supports the notion that chord roots correspond to chord tones with particularly high perceptual salience, and that this perceptual salience can be predicted by computational models of pitch detection (Parncutt, Sattmann, Gaich, & Seither-Preisler, 2019).

These empirical results motivate the use of root-finding models for simulating higher-level cognitive representations of harmonic structure. Various root-finding models exist, some deriving from psychoacoustics, some deriving from music theory, and some deriving from statistical modeling. A key example of the psychoacoustic approach is Parncutt’s (1988) root-finding model, which was derived from Terhardt’s (1982) root-finding model and subsequently extended in Parncutt (1997). The model provides a straightforward hand-computable approach to estimating the chord root for a given

pitch-class set; we have implemented it in the freely available R package *parn88*.¹⁷

Parncutt’s (1989) model relaxes the octave equivalence assumption and provides a much more detailed account of peripheral psychoacoustic processes such as masking (see also Parncutt & Strasburger, 1994); we have implemented this model in the R package *parn94*.¹⁸ Examples of the music-theoretic approach include Temperley’s (1997) algorithm, Pardo and Birmingham’s (2002) template-matching algorithm, Sapp’s (2007) stack-of-thirds algorithm, and Cambouropoulos et al.’s (2014) General Chord Type algorithm. Raphael and Stoddard’s (2004) algorithm follows the statistical approach, performing harmonic analysis using Hidden Markov Models. These root-finding algorithms provide several options for simulating root-based cognitive representations of harmonic structure.

Tonality

The notion of tonality is fundamental to Western music theory. Common-practice tonality is said to be organized around 24 possible keys, where each key is defined by a) choosing one of the 12 pitch classes to be the tonic and b) choosing one of the two diatonic modes: major and minor. The chosen key is then reflected in the distribution of notes selected by the composer. The seven most common pitch classes in a given key constitute the key’s scale, from which most of the pitch material is drawn.

Music theorists often represent chord progressions relative to the tonic. Such representations may be termed *scale-degree representations*, because each pitch class is represented as a degree of the key’s underlying scale. It is thought that Western listeners also form implicit scale-degree representations when listening to Western tonal music (e.g., Arthur, 2018). The key challenge in constructing such representations automatically is to simulate the process of *key finding*; that is, given a musical extract, to estimate the underlying key at each point in the extract.

¹⁷ <https://github.com/pmcharrison/parn88>

¹⁸ <https://github.com/pmcharrison/parn94>

Various key-finding models exist in the literature. Arguably most influential is the Krumhansl-Schmuckler algorithm (Krumhansl, 1990), which estimates the key of a given musical passage by correlating pitch-class prevalences in a given musical extract (termed *key profiles*) with experimentally determined key profiles for major and minor keys. Several variants on this approach have subsequently been presented, with these variants modifying the key profiles (Aarden, 2003; Albrecht & Shanahan, 2013; Bellmann, 2005; Sapp, 2011; Temperley, 1999), replacing key profiles with interval profiles (Madsen & Widmer, 2007), adding psychoacoustic modeling (Huron & Parncutt, 1993), implementing sliding windows (Shmulevich & Yli-Harja, 2000), and incorporating self-organizing maps (Schmuckler & Tomovski, 2005). Outside the Krumhansl-Schmuckler tradition, several algorithms instead take statistical approaches to key finding, including Temperley’s (2007) Bayesian model, Hu and Saul’s (2009) probabilistic topic model, Quinn’s (2010) chord-progression model, and White’s (2018) feedback model. Chew (2002) has also presented an approach for identifying key regions within a musical piece based on the geometric Spiral Array model (Chew, 2000). These models provide many options for simulating key-based representations of harmonic progressions.

For music theorists, tonality typically implies more than key finding: it also describes high-level principles such as how certain chords are more stable than others, how certain chords tend to resolve to other chords, and how harmonic progressions can be organized into hierarchies that potentially span the entire musical piece. These principles are typically characterized as both musical and perceptual: they are intended to describe both how tonal music is created by composers and how it is perceived by listeners. Various models have been developed to summarize these principles of tonal harmony, which typically embed chords within some kind of ‘tonal space’ within which the notion of distance between chords is well-defined. One prominent example is Lerdahl’s (1988) Tonal Pitch Space model, which operates on a symbolic level and is based on a hierarchical conceptualization of the Western chromatic scale. A contrasting example is the Tonal

Space model of Janata et al. (2002; see also Collins et al., 2014), which projects audio input onto a toroidal (‘donut-shaped’) self-organizing map. Pretrained on an artificial melody that modulates through all 24 diatonic keys, the Tonal Space model is supposed to learn a topological representation of Western tonality (see also Toivianen & Krumhansl, 2003). These kinds of models may be used to simulate tonality-based cognitive representations of harmonic progressions.

Functional harmony

Music theorists often assign chords to functional categories such as ‘tonic’, ‘dominant’, and ‘subdominant’. These categories reflect the chord’s implications for surrounding chords: for example, dominant chords are particularly likely to be followed by tonic chords. The categories are typically associated with certain chord roots, expressed relative to the prevailing tonality: for example, the dominant chord category is particularly associated with chord roots situated on the fifth degree of the prevailing diatonic scale. It is possible that untrained Western listeners also possess implicit knowledge of these functional categories, acquired through incidental exposure to Western music. Recent work has demonstrated that the acquisition of these functional categories can be simulated using computational techniques such as hidden Markov models (White & Quinn, 2018) and the information bottleneck algorithm (Jacoby, Tishby, & Tymoczko, 2015). Such algorithms could be incorporated within cognitive models of harmony perception to simulate implicit knowledge of functional chord categories.

Corpus translation

Music cognition research often takes advantage of large music corpora, for example to analyze the cognitive underpinnings of the creative process, or to simulate processes of statistical learning from lifetime musical exposure. While various corpora of chord sequences exist, these corpora are typically not expressed in the low-level representation schemes described in this paper. We will now describe various solutions to this problem.

Textual symbols

Many digitally encoded music corpora notate chords using textual symbols such as ‘min9’ and ‘sus4’ (e.g., the Billboard Corpus, Burgoyne, 2011; the iRb corpus, Broze & Shanahan, 2013). Translating these corpora into our representational network requires decoding these symbols. Unfortunately this is a poorly defined task, as the mapping between these symbols and pitch-class content is not deterministic; instead, musicians use their experience to interpret these symbols, and can exert freedom to choose different chord realizations on the basis of musical context and personal preference. Ideally, a cognitive musicologist would formalise this process by analysing many performances and quantifying the frequency with which different textual symbols are mapped to different chord pitches in different contexts. However, in the absence of such prior work, a useful stopgap is to construct a mapping between textual symbols and chord pitches on the basis of music theory, and use this mapping to translate corpora such as the Billboard Corpus and the iRb corpus into our representational network. While such a mapping will not capture all the nuances of real-life performance, it should nonetheless provide a useful approximation for supporting downstream psychoacoustic and cognitive analyses.

We have compiled such a mapping between textual symbols and chord pitches. Reflecting the typical construction of lead sheets, these textual symbols correspond to *chord qualities*, expressing the pitch-class content of a chord relative to the chord’s root. This mapping can be accessed through the function `decode_chord_quality` in the *hrep* R package. Assuming that the chord’s root and bass notes can be extracted by an automatic parser, this mapping is theoretically sufficient to convert the chord symbols of a lead sheet into pitch-class chord representations. In practice, new corpora and notation schemes are likely to contain occasional textual symbols not present in our mapping; however, these cases should be rare enough for the researcher to update the mapping manually. We would like to encourage future researchers to submit such updates to the package’s online

repository,¹⁹ so that the mapping can become more comprehensive with time. Using this mapping, we have translated the Billboard popular music corpus and the iRb corpus into pitch-class chord notation. The resulting corpora are available in our *hcorp* package.²⁰

It is often be useful to perform the reverse operation, mapping integer-based representations of chords to textual symbols. To this end, Hedges and Wiggins (2016) provide a useful algorithm that takes as input a pitch-class set, which should be expressed relative to the chord root, and returns a textual label corresponding to the chord quality, for example ‘min7’, ‘dim’, or ‘sus’. Combined with a root-finding algorithm (e.g., Cambouropoulos et al., 2014; Parncutt, 1988, 1997; Sapp, 2007), this algorithm could be used to convert the symbolic representations described here into a text-based format readable by musicians.

These corpora typically do not provide any information about chord voicing, which is unproblematic for most corpus analyses, but is problematic for generating stimuli for perceptual experiments. It is possible to address this problem using simple heuristics, such as assigning the bass pitch class to the octave below middle C and the remaining pitch classes to the octave above middle C (e.g., Harrison & Pearce, 2018b), but the outcome is typically aesthetically unsatisfying and musically unrealistic. To address this problem, we have developed a data-driven algorithm that finds optimized voicings for chord sequences on the basis of a variety of psychoacoustic features (Harrison & Pearce, 2020a). This algorithm is implemented in a publicly available R package called *voicer*.²¹

¹⁹ <http://hrep.pmcharrison.com>

²⁰ <https://github.com/pmcharrison/hcorp>

²¹ <https://github.com/pmcharrison/voicer>

Polyphonic scores

Music corpora from the tradition of Western art music typically do not notate chord sequences explicitly, but instead notate the full polyphonic texture of the original musical scores. An important challenge for harmonic analysis is then to derive chord sequences from these polyphonic textures.

One approach is *full expansion* (Conklin, 2002), also known as ‘salami slicing’ or ‘chordifying’. Here the composition is partitioned into chords at each point where a new onset occurs in any voice.²² Each chord is then defined as the full set of pitches that sound within the corresponding partition. An advantage of this approach is that it is straightforward to automate, and it captures much of the detail present in the original musical score. We have translated a corpus of 370 chorale harmonizations by J. S. Bach into pitch chords using this approach, and have made the resulting corpus available in the *hcorp* package.²³

A limitation of full expansion is that it fails to differentiate the roles of different tones within a sonority. Music theorists often differentiate tones into ‘chord tones’ and ‘non-chord tones’, where the chord tones represent a prototypical sonority such as a major or minor triad, and the non-chord tones embellish this sonority through techniques such as suspensions and passing notes. Depending on the application, a researcher may wish to exclude chord tones from their chord sequences, and potentially add implied but missing chord tones (e.g., the fifth in a seventh chord). This differentiation of chord tones from non-chord tones is vital for chord labeling and for functional harmonic analysis. Unfortunately, the process is rather subjective, as epitomized for example by debates over the analysis of the Tristan chord (Martin, 2008). Nonetheless, computational approximations to this process are possible. For example, Rohrmeier and Cross (2008)

²² Alternatively, one could partition at both onset and offset points.

²³ The original corpus was sourced from KernScores (<http://kern.humdrum.org>; Sapp, 2005).

provide a simple approach for chorale harmonizations by J. S. Bach, where pitch-class sets are generated by full expansion, and then one pitch-class set is chosen for each quarter-note segment of the chorale, with this pitch-class set being chosen to minimize dissonance.

Pardo and Birmingham (2002) present a rule-based system where chord qualities are inferred using template matching, and partition points are found using dynamic programming. Alternative approaches to chord-tone detection and chord labeling can be found in Barthélemy and Bonardi (2001), Chen and Su (2018), Chuan and Chew (2011), Ju et al. (2017), Kröger et al. (2008), Masada and Bunescu (2017), Mearns (2013), Raphael and Stoddard (2004), Temperley and Sleator (1999), and Winograd (1968). These works vary in the extent to which they apply explicit music-theoretic rules; recent work minimizes the role of explicit knowledge, and instead uses algorithms that learn these principles from musical data. In almost all cases, these approaches are targeted towards replicating the annotations of expert music theorists; we are unaware of any such work actively engaging with empirical studies of music perception. This strategy is perfectly acceptable for many applications in music informatics, but it is dissatisfying for applications in music cognition research. Future work should ideally examine the psychological foundations of chord-tone identification and chord labeling and use this knowledge to improve the cognitive validity of these algorithms.

Audio

Chord sequences can also be derived from audio files. This is generally a harder task than deriving chord sequences from symbolic music representations, but one that has received much attention in recent years. Effective chord transcription algorithms could vastly widen the amount of data available for music corpus analyses.

Many systems have been proposed for deriving sequences of chord labels from audio files (e.g., Boulanger-Lewandowski, Bengio, & Vincent, 2013; Haas, Magalhães, & Wiering, 2012; Lee & Slaney, 2008; Mauch, 2010; Mauch, Noland, & Dixon, 2009). These systems

are typically data-driven, training either Hidden Markov Models or deep neural networks on datasets pairing audio frames with chord labels, and producing algorithms capable of generating chord labels for unseen audio. These chord labels could then be translated to our representational network using the methods described above.

One could alternatively begin by automatically transcribing the full polyphonic score. Automatic polyphonic transcription remains a very difficult task, but substantial progress has been made in recent years (e.g., Benetos & Dixon, 2013; Boulanger-Lewandowski, Vincent, & Bengio, 2012; Sigtia, Benetos, & Dixon, 2016). Having transcribed a full polyphonic score, chord sequences could then be derived using full expansion, potentially also applying a symbolic chord-tone detection algorithm.

Case studies

We will now describe two case studies illustrating how the present framework has proved useful in our own research; hopefully these examples will suggest future worthwhile applications to the reader.

Consonance

When we first started exploring the psychology of musical consonance, we were pleased to see that so many consonance models existed in the literature, but disappointed to realize that many of these models lacked systematic comparative evaluations or public implementations. We therefore set about implementing and evaluating a large number of these models, and on this basis constructing a composite model that would predict an input chord’s consonance as a weighted sum of several sub-models capturing different psychological mechanisms (Harrison & Pearce, 2020c). We found that, in practice, a significant amount of code complexity could be saved by consolidating all the representational considerations in one codebase. Our resulting analyses represented each input chord as:

- a) a smooth pitch-class spectrum for Milne’s (2013) and Harrison and Pearce’s (2018a) harmonicity models;
- b) a sparse pitch spectrum for four of the interference models (Hutchinson & Knopoff, 1978; Sethares, 1993; Vassilakis, 2001; Wang, Shen, Guo, Tang, & Hamade, 2013);
- c) a pitch chord type for Stolzenburg’s (2015) and Gill and Purves’ (2009) harmonicity models;
- d) a pitch-class set for Johnson-Laird et al.’s (2012) model;
- e) a sensory variant of the sparse pitch spectrum for Parncutt and Strasburger’s (1994) models;
- f) a pitch-class chord type for our new cultural model;
- g) audio for the Essentia (Bogdanov et al., 2013) and MIRtoolbox (Lartillot et al., 2008) packages.

Our *hrep* package provided a central solution for these representation manipulations and a coherent terminology for describing them, significantly enhancing the efficiency and scalability of the project.

Harmonic syntax

In the second project we were interested in analyzing how chord transitions in different musical styles might be explained by psychoacoustic models (Harrison & Pearce, 2018a). We wished to avoid *a priori* assumptions about Western tonality, so we decided to use the pitch-class set as the basic representation, and we analyzed the statistics of pairwise transitions between pitch-class sets, aiming to explain them in terms of four features:

- a) the voice-leading distance between the two pitch-class sets,
- b) the similarity of the two sonorities’ smooth pitch-class spectra,
- c) the harmonicity of the target sonority’s smooth pitch-class spectrum;
- d) the cardinality of the target pitch-class set.

There are 4,095 possible pitch-class sets in 12-tone equal temperament, and therefore 16,769,025 possible transitions between these pitch-class sets, which means a lot of transition probabilities to model. However, we had selected each of our features to be transposition-invariant, thereby allowing for some significant computational simplifications. Specifically, we transposed each of our tabulated chord transitions such that the first pitch-class set was replaced by its pitch-class set type, thereby reducing the number of potential starting pitch-class sets to 351, and correspondingly reducing the total number of possible transitions ten-fold. The encoding system described in Appendix B provided a natural way to enumerate these different sonorities in transition matrices. Taking advantage of these invariances and efficient integer encodings significantly reduced the computational demands of the algorithm, and made it practical to add a computationally expensive bootstrapping procedure.

Conclusion

This paper’s primary contribution is cataloguing a collection of low-level harmonic representations for music cognition research. Our description clarifies the implicit cognitive assumptions of each representation, and provides computational methods for translating between representations. We hope that this explicit account will assist cognitive scientists by providing a terminological reference for these representations, providing guidance on choosing representations, and helping researchers to reconcile music corpora from heterogeneous data sources.

A second contribution is the creation of the open-source R package *hrep*, which implements these different representations and encodings in an easy-to-use object-oriented framework. In our experience, a large part of harmony modeling projects tends to be occupied with the scientifically uninteresting task of taking musical corpora and translating them to the appropriate representation for a given statistical model or feature extractor. By centralizing these processes in a single R package, we expect that we can save

researchers significant time in this research pipeline, and facilitate rapid prototyping of different analysis approaches. Of course, many music researchers use other programming languages such as Python and Lisp, and the *hrep* package will be less useful here. However, the *hrep* package can still be used for preprocessing musical inputs to programs written in these languages, and it provides a template according to which analogous harmonic representation packages could be developed.

We have discussed how higher-level cognitive representations, such as voice leadings, chord roots, tonality, and functional harmony, can be derived from the representations included here. These representations are typically more computationally complex to implement, and are still active topics of research, so we have not implemented them in our *hrep* package. However, we have reviewed the various computational methods for deriving these higher-level representations from the low-level representations described here, and have made some of these methods available in standalone R packages.

We also discussed how to translate different musical corpora into these low-level representations. Various approaches exist, depending on whether the corpus is represented as textual chord symbols, polyphonic scores, or audio. One contribution of this paper is to provide a systematic mapping between common chord symbols and the pitch classes implied by these chord symbols. This should help future researchers translate music corpora into the representations described here.

In our eyes, the main limitation of our approach is assuming that the input data are already represented as sequences of block chords. Unfortunately, real music does not always have such a simple structure. Chords are often formed implicitly by the superposition of multiple melodic or contrapuntal parts, and elaborated by structures such as passing notes and appoggiaturas. An important goal for future cognitive research is to examine the way in which listeners extract chord-like representations from these complex musical textures.

References

- Aarden, B. J. (2003). *Dynamic melodic expectancy* (PhD thesis). Ohio State University, Columbus, OH.
- Albrecht, J., & Shanahan, D. (2013). The use of large corpora to train a new type of key-finding algorithm: An improved treatment of the minor mode. *Music Perception*, 31, 59–67. <https://doi.org/10.1525/MP.2013.31.1.59>
- Arthur, C. (2018). A perceptual study of scale-degree qualia in context. *Music Perception*, 35, 295–314. <https://doi.org/10.1525/MP.2018.35.3.295>
- Babbitt, M. (1965). The use of computers in musicological research. *Perspectives of New Music*, 3(2), 74–83.
- Bachem, A. (1950). Tone height and tone chroma as two different pitch qualities. *Acta Psychologica*, 7, 80–88. [https://doi.org/10.1016/0001-6918\(50\)90004-7](https://doi.org/10.1016/0001-6918(50)90004-7)
- Barthélemy, J., & Bonardi, A. (2001). Figured bass and tonality recognition. In *Proceedings of the Second International Symposium on Music Information Retrieval* (pp. 129–136). Bloomington, IN.
- Bellmann, H. (2005). About the determination of key of a musical excerpt. In R. Kronland-Martinet, T. Voinier, & S. Ystad (Eds.), *Proceedings of the Third International Symposium of Computer Music Modeling and Retrieval* (pp. 76–91). Berlin, Germany: Springer.
- Benetos, E., & Dixon, S. (2013). Multiple-instrument polyphonic music transcription using a temporally constrained shift-invariant model. *The Journal of the Acoustical Society of America*, 133, 1727–1741. <https://doi.org/10.1121/1.4790351>
- Bogdanov, D., Wack, N., Gómez, E., Gulati, S., Herrera, P., Mayor, O., ... Serra, X. (2013).

- Essentia: An audio analysis library for music information retrieval. In *Proceedings of the 14th International Society for Music Information Retrieval Conference*. Curitiba, Brazil.
- Boulanger-Lewandowski, N., Bengio, Y., & Vincent, P. (2013). Audio chord recognition with recurrent neural networks. In *Proceedings of the 14th International Society for Music Information Retrieval Conference* (pp. 335–340). Curitiba, Brazil.
- Boulanger-Lewandowski, N., Vincent, P., & Bengio, Y. (2012). Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. In J. Langford & J. Pineau (Eds.), *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*. Edinburgh, Scotland. Retrieved from <http://arxiv.org/abs/1206.6392>
- Bregman, A. S. (1990). *Auditory scene analysis: The perceptual organization of sound*. Cambridge, MA: MIT Press.
- Broze, Y., & Shanahan, D. (2013). Diachronic changes in jazz harmony: A cognitive perspective. *Music Perception*, 31, 32–45. <https://doi.org/10.1525/rep.2008.104.1.92>
- Burgoyne, J. A. (2011). *Stochastic Processes & Database-Driven Musicology* (PhD thesis). McGill University, Montréal, Québec, Canada. <https://doi.org/10.1145/957013.957041>
- Callender, C., Quinn, I., & Tymoczko, D. (2008). Generalized voice-leading. *Science*, 320, 346–348. <https://doi.org/10.1126/science.1153021>
- Cambouropoulos, E. (2016). The harmonic musical surface and two novel chord representation schemes. In D. Meredith (Ed.), *Computational music analysis* (pp. 31–56). Cham, Switzerland: Springer International Publishing.
- Cambouropoulos, E., Kaliakatsos-Papakostas, M., & Tsougras, C. (2014). An idiom-independent representation of chords for computational music analysis and

- generation. In *Proceedings of the joint 11th Sound and Music Computing Conference (SMC) and 40th International Computer Music Conference (ICMC)*. Athens, Greece.
- Carter, E. (2002). *Harmony Book*. (N. Hopkins & J. F. Link, Eds.). New York, NY: Carl Fischer.
- Chen, T.-P., & Su, L. (2018). Functional harmony recognition of symbolic music data with multi-task recurrent neural networks. In E. Gómez, Hu Xiao, E. Humphrey, & E. Benetos (Eds.), *Proceedings of the 19th International Society for Music Information Retrieval Conference* (pp. 90–97). Paris, France.
- Cheung, V. K. M., Harrison, P. M. C., Meyer, L., Pearce, M. T., Haynes, J.-D., & Koelsch, S. (2019). Uncertainty and surprise jointly predict musical pleasure and amygdala, hippocampus, and auditory cortex activity. *Current Biology*, 29, 4084–4092.
<https://doi.org/10.1016/j.cub.2019.09.067>
- Chew, E. (2000). *Towards a mathematical model of tonality* (PhD thesis). MIT, Cambridge, MA.
- Chew, E. (2002). The Spiral Array: An algorithm for determining key boundaries. In C. Anagnostopoulou, M. Ferrand, & A. Smaill (Eds.), *Music and Artificial Intelligence. ICMAI 2002. Lecture Notes in Computer Science, vol. 2445*. Berlin, Germany: Springer. <https://doi.org/10.1007/3-540-45722-4>
- Chuan, C.-H., & Chew, E. (2011). Generating and evaluating musical harmonizations that emulate style. *Computer Music Journal*, 35(4), 64–82.
- Collins, T., Tillmann, B., Barrett, F. S., Delbé, C., & Janata, P. (2014). A combined model of sensory and cognitive representations underlying tonal expectations in music: From audio signals to behavior. *Psychological Review*, 121, 33–65.
<https://doi.org/10.1037/a0034695>

- Conklin, D. (2002). Representation and discovery of vertical patterns in music. In C. Anagnostopoulou, M. Ferrand, & A. Smaill (Eds.), *Music and Artificial Intelligence: Proc. ICMAI 2002* (pp. 32–42). Berlin, Germany: Springer-Verlag.
- Dean, R. T., Milne, A. J., & Bailes, F. (2019). Spectral pitch similarity is a predictor of perceived change in sound- as well as note-based music. *Music & Science*, 2. <https://doi.org/10.1177/2059204319847351>
- Dienes, Z., & Longuet-Higgins, C. (2004). Can musical transformations be implicitly learned? *Cognitive Science*, 28, 531–558. <https://doi.org/10.1016/j.cogsci.2004.03.003>
- Dillon, G. (2013). Calculating the dissonance of a chord according to Helmholtz. *European Physical Journal Plus*, 128(90). <https://doi.org/10.1140/epjp/i2013-13090-4>
- Forte, A. (1973). *The structure of atonal music*. New Haven, CT: Yale University Press.
- Gill, K. Z., & Purves, D. (2009). A biological rationale for musical scales. *PLoS ONE*, 4(12). <https://doi.org/10.1371/journal.pone.0008144>
- Haas, B. de, Magalhães, J. P., & Wiering, F. (2012). Improving audio chord transcription by exploiting harmonic and metric knowledge. In A. de Souza Britto Jr., F. Gouyon, & S. Dixon (Eds.), *Proceedings of the 13th International Society for Music Information Retrieval Conference* (pp. 295–300). Porto, Portugal.
- Harrison, P. M. C., & Pearce, M. T. (2018a). An energy-based generative sequence model for testing sensory theories of Western harmony. In E. Gómez, Hu Xiao, E. Humphrey, & E. Benetos (Eds.), *Proceedings of the 19th International Society for Music Information Retrieval Conference*. Paris, France.
- Harrison, P. M. C., & Pearce, M. T. (2018b). Dissociating sensory and cognitive theories of harmony perception through computational modeling. In R. Parncutt & S. Sattmann

- (Eds.), *Proceedings of ICMPC15/ESCOM10*. Graz, Austria.
<https://doi.org/10.31234/osf.io/wgjyv>
- Harrison, P. M. C., & Pearce, M. T. (2020a). A computational cognitive model for the analysis and generation of chord voicings. *Music Perception*, *37*, 208–224.
<https://doi.org/10.1525/mp.2020.37.3.208>
- Harrison, P. M. C., & Pearce, M. T. (2020b). A computational cognitive model for the analysis and generation of voice leadings. *Music Perception*, *37*, 208–224.
- Harrison, P. M. C., & Pearce, M. T. (2020c). Simultaneous consonance in music perception and composition. *Psychological Review*, *127*, 216–244.
<https://doi.org/10.1037/rev0000169>
- Harte, C., Sandler, M., Abdallah, S., & Gómez, E. (2005). Symbolic representation of musical chords: A proposed syntax for text annotations. In *Proceedings of the 6th International Conference on Music Information Retrieval* (pp. 66–71). London, UK.
- Hedges, T., & Wiggins, G. A. (2016). The prediction of merged attributes with multiple viewpoint systems. *Journal of New Music Research*, *45*, 314–332.
<https://doi.org/10.1080/09298215.2016.1205632>
- Hu, D. J., & Saul, L. K. (2009). A probabilistic topic model for unsupervised learning of musical key-profiles. In K. Hirata, G. Tzanetakis, & K. Yoshii (Eds.), *Proceedings of the 10th International Society for Music Information Retrieval Conference*. Kobe, Japan.
- Huron, D., & Parncutt, R. (1993). An improved model of tonality perception incorporating pitch salience and echoic memory. *Psychomusicology*, *12*, 154–171.
- Hutchinson, W., & Knopoff, L. (1978). The acoustic component of Western consonance. *Journal of New Music Research*, *7*, 1–29. <https://doi.org/10.1080/09298217808570246>

- Immerseel, L. V., & Martens, J. (1992). Pitch and voiced/unvoiced determination with an auditory model. *The Journal of the Acoustical Society of America*, *91*, 3511–3526.
- Jacoby, N., Tishby, N., & Tymoczko, D. (2015). An information theoretic approach to chord categorization and functional harmony. *Journal of New Music Research*, *44*, 219–244. <https://doi.org/10.1080/09298215.2015.1036888>
- Jacoby, N., Undurraga, E. A., McPherson, M. J., Valdés, J., Ossandón, T., & McDermott, J. H. (2019). Universal and non-universal features of musical pitch perception revealed by singing. *Current Biology*, *29*(19), 3229–3243.
- Janata, P., Birk, J. L., Van Horn, J. D., Leman, M., Tillmann, B., & Bharucha, J. J. (2002). The cortical topography of tonal structures underlying Western music. *Science*, *298*, 2167–2170. <https://doi.org/10.1126/science.1076262>
- Johnson-Laird, P. N., Kang, O. E., & Leong, Y. C. (2012). On musical dissonance. *Music Perception*, *30*, 19–35.
- Ju, Y., Condit-Schultz, N., Arthur, C., & Fujinaga, I. (2017). Non-chord tone identification using deep neural networks. In K. Page (Ed.), *Proceedings of the 4th International Workshop on Digital Libraries for Musicology* (pp. 13–16). Shanghai, China. <https://doi.org/10.1145/3144749.3144753>
- Kröger, P., Passos, A., Sampaio, M. da S., & Cidra, G. de. (2008). Rameau: A system for automatic harmonic analysis. In *Proceedings of the 2008 International Computer Music Conference* (pp. 273–281). Belfast, Northern Ireland.
- Krumhansl, C. L. (1990). *Cognitive foundations of musical pitch*. New York, NY: Oxford University Press.
- Krumhansl, C. L., Sandell, G. J., & Sergeant, D. C. (1987). The perception of tone

- hierarchies and mirror forms in twelve-tone serial music. *Music Perception*, 5, 31–77.
- Lartillot, O., Toivianen, P., & Eerola, T. (2008). A Matlab toolbox for Music Information Retrieval. In C. Preisach, H. Burkhardt, L. Schmidt-Thieme, & R. Decker (Eds.), *Data analysis, machine learning and applications* (pp. 261–268). Berlin, Germany: Springer.
- Lee, K., & Slaney, M. (2008). Acoustic chord transcription and key extraction from audio using key-dependent HMMs trained on synthesized audio. *IEEE Transactions on Audio, Speech and Language Processing*, 16, 291–301.
<https://doi.org/10.1109/TASL.2007.914399>
- Leman, M. (2000). An auditory model of the role of short-term memory in probe-tone ratings. *Music Perception*, 17, 481–509.
- Lerdahl, F. (1988). Tonal pitch space. *Music Perception*, 5, 315–349.
- Lewin, D. (1987). *Generalized musical intervals and transformations*. New Haven, CT: Yale University Press.
- Madsen, S. T., & Widmer, G. (2007). Key-finding with interval profiles. In *Proceedings of the International Computer Music Conference (ICMC)*. Copenhagen, Denmark.
- Martin, N. (2008). The Tristan chord resolved. *Intersections*, 28(2), 6–30.
<https://doi.org/10.7202/029953ar>
- Martino, D. (1961). The source set and its aggregate formations. *Journal of Music Theory*, 5, 224–273.
- Masada, K., & Bunescu, R. (2017). Chord recognition in symbolic music using semi-Markov Conditional Random Fields. In S. J. Cunningham, Z. Duan, X. Hu, & D. Turnbull (Eds.), *Proceedings of the 18th International Society for Music Information Retrieval Conference* (pp. 272–278). Suzhou, China.

- Mauch, M. (2010). *Automatic chord transcription from audio using computational models of musical context* (PhD thesis). Queen Mary University of London, London, UK.
- Mauch, M., & Dixon, S. (2008). A discrete mixture model for chord labelling. In *Proceedings of the 9th International Society for Music Information Retrieval Conference* (pp. 45–50). Philadelphia, PA.
- Mauch, M., Dixon, S., Harte, C., Casey, M., & Fields, B. (2007). Discovering chord idioms through Beatles and Real Book songs. In *Proceedings of the 8th International Society for Music Information Retrieval Conference*. Vienna, Austria.
- Mauch, M., Noland, K., & Dixon, S. (2009). Using musical structure to enhance automatic chord transcription. In K. Hirata, G. Tzanetakis, & K. Yoshii (Eds.), *Proceedings of the 10th International Society for Music Information Retrieval Conference* (pp. 231–236). Kobe, Japan.
- Mearns, L. (2013). *The computational analysis of harmony in Western art music* (PhD thesis). Queen Mary University of London, London, UK.
- Milne, A. J. (2013). *A computational model of the cognition of tonality* (PhD thesis). The Open University, Milton Keynes, England.
- Milne, A. J., & Holland, S. (2016). Empirically testing Tonnetz, voice-leading, and spectral models of perceived triadic distance. *Journal of Mathematics and Music*, 10, 59–85.
<https://doi.org/10.1080/17459737.2016.1152517>
- Milne, A. J., Laney, R., & Sharp, D. (2015). A spectral pitch class model of the probe tone data and scalar tonality. *Music Perception*, 32, 364–393.
<https://doi.org/10.1525/MP.2015.32.4.364>
- Milne, A. J., Laney, R., & Sharp, D. B. (2016). Testing a spectral model of tonal affinity

- with microtonal melodies and inharmonic spectra. *Musicae Scientiae*, 20, 465–494.
<https://doi.org/10.1177/1029864915622682>
- Milne, A. J., Sethares, W. A., Laney, R., & Sharp, D. B. (2011). Modelling the similarity of pitch collections with expectation tensors. *Journal of Mathematics and Music*, 5, 1–20. <https://doi.org/10.1080/17459737.2011.573678>
- Moore, B. C. J., Peters, R. W., & Glasberg, B. R. (1985). Thresholds for the detection of inharmonicity in complex tones. *The Journal of the Acoustical Society of America*, 77(5), 1861–1867.
- Moss, F. C., Neuwirth, M., Harasim, D., & Rohrmeier, M. (2019). Statistical characteristics of tonal harmony: A corpus study of Beethoven’s string quartets. *PLoS One*, 14(6). <https://doi.org/10.1371/journal.pone.0217242>
- Neuwirth, M., Harasim, D., Moss, F. C., & Rohrmeier, M. A. (2018). The Annotated Beethoven Corpus (ABC): A dataset of harmonic analyses of all Beethoven string quartets. *Frontiers in Digital Humanities*, 5. <https://doi.org/10.3389/fdigh.2018.00016>
- Pardo, B., & Birmingham, W. P. (2002). Algorithms for chordal analysis. *Computer Music Journal*, 26(2), 27–49. <https://doi.org/10.1162/014892602760137167>
- Parncutt, R. (1988). Revision of Terhardt’s psychoacoustical model of the root(s) of a musical chord. *Music Perception*, 6, 65–94.
- Parncutt, R. (1989). *Harmony: A psychoacoustical approach*. Berlin, Germany: Springer-Verlag.
- Parncutt, R. (1997). A model of the perceptual root(s) of a chord accounting for voicing and prevailing tonality. In M. Leman (Ed.), *Music, gestalt, and computing: Studies in cognitive and systematic musicology* (pp. 181–199). Berlin, Germany: Springer.

<https://doi.org/10.1007/BFb0034114>

- Parncutt, R., Sattmann, S., Gaich, A., & Seither-Preisler, A. (2019). Tone profiles of isolated musical chords: Psychoacoustic versus cognitive models. *Music Perception*, 36, 406–430.
- Parncutt, R., & Strasburger, H. (1994). Applying psychoacoustics in composition: “Harmonic” progressions of “nonharmonic” sonorities. *Perspectives of New Music*, 32(2), 88–129.
- Plantinga, J., & Trainor, L. J. (2005). Memory for melody: Infants use a relative pitch code. *Cognition*, 98, 1–11. <https://doi.org/10.1016/j.cognition.2004.09.008>
- Plomp, R., & Levelt, W. J. M. (1965). Tonal consonance and critical bandwidth. *The Journal of the Acoustical Society of America*, 38, 548–560.
<https://doi.org/10.1121/1.1909741>
- Quinn, I. (2010). Are pitch-class profiles really “key for key”? *Zeitschrift Der Gesellschaft Für Musiktheorie*, 7(2), 151–163.
- Quinn, I., & Mavromatis, P. (2011). Voice-leading prototypes and harmonic function. In C. Agón, M. Andreatta, G. Assayag, E. Amiot, J. Bresson, & J. Mandereau (Eds.), *Mathematics and Computation in Music – Third International Conference, MCM 2011* (pp. 230–240). Berlin, Germany: Springer.
- Rahn, J. (1980). *Basic atonal theory*. New York, NY: Schirmer.
- Rameau, J.-P. (1722). *Treatise on harmony*. Paris, France: Jean-Baptiste-Christophe Ballard.
- Raphael, C., & Stoddard, J. (2004). Functional harmonic analysis using probabilistic models. *Computer Music Journal*, 28(3), 45–52.

<https://doi.org/10.1162/0148926041790676>

- Rohrmeier, M. A., & Cross, I. (2008). Statistical properties of tonal harmony in Bach's chorales. In K. Miyazaki, Y. Hiraga, M. Adachi, Y. Nakajima, & M. Tsuzaki (Eds.), *Proceedings of the 10th International Conference on Music Perception and Cognition* (pp. 619–627). Sapporo, Japan.
- Sapp, C. (2007). Computational chord-root identification in symbolic musical data: Rationale, methods, and applications. *Computing in Musicology*, 15, 99–119.
- Sapp, C. S. (2005). Online database of scores in the Humdrum file format. In *Proceedings of the 6th International Society for Music Information Retrieval Conference* (pp. 664–665).
- Sapp, C. S. (2011). *Computational methods for the analysis of musical structure* (PhD thesis). Stanford University, Stanford, CA.
- Schmuckler, M. A., & Tomovski, R. (2005). Perceptual tests of an algorithm for musical key-finding. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 1124–1149. <https://doi.org/10.1037/0096-1523.31.5.1124>
- Sears, D. R. W., Arzt, A., Frostel, H., Sonnleitner, R., & Widmer, G. (2017). Modeling harmony with skip-grams. In S. J. Cunningham, Z. Duan, X. Hu, & D. Turnbull (Eds.), *Proceedings of the 18th International Society for Music Information Retrieval Conference*. Suzhou, China.
- Sethares, W. A. (1993). Local consonance and the relationship between timbre and scale. *The Journal of the Acoustical Society of America*, 94, 1218–1228.
- Sethares, W. A. (2005). *Tuning, timbre, spectrum, scale*. London, UK: Springer.
- Shmulevich, I., & Yli-Harja, O. (2000). Localized key finding: Algorithms and applications.

Music Perception, 17, 531–544.

Sigtia, S., Benetos, E., & Dixon, S. (2016). An end-to-end neural network for polyphonic music transcription. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24, 927–939.

Stainsby, T., & Cross, I. (2009). The perception of pitch. In S. Hallam, I. Cross, & M. Thaut (Eds.), *The Oxford handbook of music psychology* (pp. 47–58). New York, NY: Oxford University Press.

Stolzenburg, F. (2015). Harmony perception by periodicity detection. *Journal of Mathematics and Music*, 9, 215–238. <https://doi.org/10.1080/17459737.2015.1033024>

Temperley, D. (1997). An algorithm for harmonic analysis. *Music Perception*, 15, 31–68.

Temperley, D. (1999). What's key for key? The Krumhansl-Schmuckler key-finding algorithm reconsidered. *Music Perception*, 17, 65–100.

Temperley, D. (2007). *Music and Probability*. Cambridge, MA: MIT Press.

Temperley, D., & De Clercq, T. (2013). Statistical analysis of harmony and melody in rock music. *Journal of New Music Research*, 423, 187–204.
<https://doi.org/10.1080/09298215.2013.788039>

Temperley, D., & Sleator, D. (1999). Modeling meter and harmony: A preference-rule approach. *Computer Music Journal*, 23(1), 10–27.

Terhardt, E. (1974). Pitch, consonance, and harmony. *The Journal of the Acoustical Society of America*, 55, 1061–1069.

Terhardt, E. (1982). Die psychoakustischen Grundlagen der musikalischen Akkordgrundtöne und deren algorithmische Bestimmung. In C. Dahlhaus & M. Krause

- (Eds.), *Tiefenstruktur*. Berlin, Germany: Technical University of Berlin.
- Terhardt, E., Stoll, G., & Seewann, M. (1982). Algorithm for extraction of pitch and pitch salience from complex tonal signals. *The Journal of the Acoustical Society of America*, 71, 679–688. <https://doi.org/10.1121/1.387544>
- Thagard, P. (2019). Cognitive science. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2019). Stanford, CA: Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/spr2019/entries/cognitive-science>
- Toiviainen, P., & Krumhansl, C. L. (2003). Measuring and modeling real-time responses to music: The dynamics of tonality induction. *Perception*, 32, 741–766. <https://doi.org/10.1068/p3312>
- Tymoczko, D. (2006). The geometry of musical chords. *Science*, 313, 72–74. <https://doi.org/10.1126/science.1126287>
- Tymoczko, D. (2008). Scale theory, serial theory and voice leading. *Music Analysis*, 27, 1–49. <https://doi.org/10.1111/j.1468-2249.2008.00257.x>
- Tymoczko, D. (2011). *A Geometry of Music*. New York, NY: Oxford University Press.
- Vassilakis, P. N. (2001). *Perceptual and physical properties of amplitude fluctuation and their musical significance* (PhD thesis). UCLA, Los Angeles, CA.
- Wang, Y. S., Shen, G. Q., Guo, H., Tang, X. L., & Hamade, T. (2013). Roughness modelling based on human auditory perception for sound quality evaluation of vehicle interior noise. *Journal of Sound and Vibration*, 332(16), 3893–3904. <https://doi.org/10.1016/j.jsv.2013.02.030>
- White, C. W. (2018). Feedback and feedforward models of musical key. *Music Theory*

Online, 24(2). <https://doi.org/10.30535/mto.24.2.4>

White, C. W., & Quinn, I. (2018). Chord context and harmonic function in tonal music.

Music Theory Spectrum, 40, 314–335O. <https://doi.org/10.1093/mts/mty021>

Winograd, T. (1968). Linguistics and the computer analysis of tonal harmony. *Journal of*

Music Theory, 12, 2–49.

Appendix A

Rahn's (1980) algorithm finds the 'normal order' (also known as the 'normal form') of a pitch-class set.²⁴ The special property of this algorithm is that it deterministically finds an ordering of the pitch-class set that is identical for all transpositions of that pitch-class set. It comprises the following steps:

1. Write the pitch-class set as an ascending list of integers.
2. Enumerate all possible 'cycles' of this list by repeatedly moving the first element to the end of the list. For example, the pitch-class set $\{0, 7, 9\}$ has three cycles: $(0, 7, 9)$, $(7, 9, 0)$, and $(9, 0, 7)$.
3. Look for the most 'compact' cycle, that is, the cycle with the smallest ascending pitch-class interval between its first and last elements. In the case of a tie, look for the cycle with the smallest ascending pitch-class interval between the first and the second-to-last element. In the case of a further tie, look at the third-to-last element, and so on. If there is still a tie, choose the cycle with the smallest initial pitch-class number. In the example above, the most compact cycle is $(7, 9, 0)$.

References

- Forte, A. (1973). *The structure of atonal music*. New Haven, CT: Yale University Press.
- Rahn, J. (1980). *Basic atonal theory*. New York, NY: Schirmer.
- Straus, J. N. (1991). A primer for atonal set theory. *College Music Symposium*, 31, 1–26.

²⁴ See Forte (1973) for a similar algorithm, and see Straus (1991) for an alternative presentation of Rahn's algorithm.

Appendix B

Here we describe algorithms for encoding four of the symbolic representations into integer form. Each algorithm assumes that the input chord is quantized to the 12-tone scale.

Pitch-class set

As defined above, a pitch-class set is a non-empty unordered set of integers between 0 and 11 inclusive (e.g. $\{0, 4, 7\}$). A pitch-class set may be encoded as follows:

1. Represent the pitch-class set as a 0-indexed binary vector of length 12, where the i th element is 1 if the pitch-class set contains the pitch class i and 0 otherwise. The pitch-class set $\{0, 4, 7\}$ would therefore receive the binary vector representation ‘100010010000’.
2. Reinterpret this binary vector as a base-2 number.
3. Convert this base-2 number to a base-10 integer.²⁵

The result is an integer encoding that ranges between 1 and 4,095; for example, $\{11\}$ is coded as 1, $\{10\}$ is coded as 2, $\{10, 11\}$ is coded as 3, and $\{0, 4, 7\}$ is coded as 2,192. The full alphabet can then be enumerated by iterating through the integers from 1 to 4,095.

Pitch-class set type

As defined above, pitch-class set types define transpositionally invariant equivalence classes over pitch-class sets; for example, the pitch-class set type $\{0, 4, 8\}$ contains the pitch-class sets $\{0, 4, 8\}$, $\{1, 5, 9\}$, $\{2, 6, 10\}$, and $\{3, 7, 11\}$. The alphabet of pitch-class set types is constructed as follows:

²⁵ A binary vector of length 12 can be converted to a base-10 integer by multiplying each i th element ($1 \leq i \leq 12$) by 2^{12-i} and summing the result.

1. Initialise an empty list.
2. Iterate sequentially through the pitch-class set alphabet, and:
 - (a) Compute the pitch-class set type of the pitch-class set;
 - (b) If this is the first occurrence of that pitch-class set type, append it to the list.

The resulting alphabet maps pitch-class set types to integers between 1 and 351; for example, $\{0\}$ maps to 1, $\{0, 1\}$ maps to 2, $\{0, 2\}$ maps to 3, and $\{0, 1, 2\}$ maps to 4.

Pitch-class chord type

As defined above, pitch-class chord types define transpositionally invariant equivalence classes over pitch-class chords; for example, the pitch-class chords $(0, \{4, 7\})$ and $(1, \{5, 8\})$ both belong to the pitch-class chord type $\{0, 4, 7\}$. By definition, the integer-vector representation of the pitch-class chord type must begin with 0; the remaining elements may then be drawn arbitrarily from the integers 1 to 11. Correspondingly, pitch-class chord types are encoded as follows:

1. Construct a 1-indexed binary vector of length 11, where the i th element is 1 if the integer vector contains i and 0 otherwise.
2. Convert the resulting binary vector to base 10.
3. Add 1.

The resulting alphabet maps pitch-class chord types to integers between 1 and 2,048; for example, $\{0\}$ maps to 1, $\{0, 1\}$ maps to 1,025, and $\{0, 4, 7\}$ maps to 145.

Pitch-class chord

Each of the 2,048 pitch-class chord types is an equivalence class containing exactly 12 pitch-class chords with 12 different bass pitch classes. Correspondingly, pitch-class chords

are encoded as follows:

1. Take the bass pitch class of the pitch-class chord, expressed as an integer.
2. Multiply it by 2,048.
3. Add the integer encoding of the pitch-class chord type.

The resulting alphabet maps pitch-class chords to integers between 1 and 24,576; for example, $(0, \{3, 6\})$ maps to 289, $(0, \{4, 7\})$ maps to 145, and $(4, \{0, 7\})$ maps to 8,457.