

Learning about harmony with Harmony Space: an overview

© Simon Holland 1994
Department of Computing,
Faculty of Mathematics, Open University, Milton Keynes, UK
Tel: UK+44 908 65 3148 email: s.holland@open.ac.uk

A version of this paper appeared as:-

Holland, S. (1994) Learning about harmony with Harmony Space: an overview. In Smith, M. and Wiggins, G. (Eds.) *Music Education: An Artificial Intelligence Approach*. Springer Verlag, London.

This paper reports on aspects of work carried out over a number of years at the Open University and the University of Aberdeen. This version was prepared at CCRMA while a visiting scholar in September 1994.

Learning about harmony with Harmony Space: an overview

© Simon Holland

Department of Computing,

Faculty of Mathematics, Open University, Milton Keynes, UK

Tel: UK-44-908-653148 email: s.holland@open.ac.uk

Abstract

Recent developments are presented in the evolution of Harmony Space, a highly enabling interface designed to encourage and facilitate rapid learning, especially by beginners, about the practical use and theory of tonal harmony, especially as applied to composition and analysis. The interface exploits both cognitive theories of tonal harmony and principles of human-computer interaction. In particular, the design of the interface draws on Balzano's and Longuet-Higgins' theories of tonal harmony. The interface allows entities of interest (notes, chords, chord progressions, key areas, modulations) to be manipulated via direct manipulation techniques using a single principled spatial metaphor. This makes it possible for novices to learn to perform a wide range of musical tasks involving harmony relatively quickly. The interface can also be used by experienced musicians to gain new insights and perform certain tasks much more easily than with conventional tools and notations. The interface is highly interactive and multi-modal, using two pointing devices and spatial, aural and kinaesthetic cues that all map uniformly into the underlying representation. Some recent implementations of Harmony Space are discussed, together with some of the musical tasks which they make tractable for beginners and experienced musicians. Aspects of the simple, consistent, principled framework behind the interface are outlined.

1 Introduction

Harmony Space is a highly interactive interface for exploring harmony and a means of learning about certain aspects of music composition. The interface draws on two cognitive theories of harmony. It exploits artificial intelligence and human computer interaction methodologies to help beginners learn about and make use of tonal harmony. The interface was originally inspired by Longuet-Higgins (1962) theory of the perception of harmony, and some current versions of the interface still strongly reflect this influence. However, for many purposes, it has turned out to be more appropriate to use Balzano's 1982 group theoretic characterisation of the perception of harmony as a grounding for the interface.

In this paper we will focus on the use of the interface for tonal and modal harmony: most versions of this interface (constructed during a series of research projects) employ a 12-fold equal-tempered division of the octave, though other implemented versions employ just intonation (Holland 1989) in order to facilitate the exploration of harmonic ideas and aspects of modulation using this tuning. Versions of the interface have also been constructed for 6-fold, 20-fold 30-fold, 42-fold 56-fold, and 72-fold divisions of the octave, to facilitate the exploration of microtonal scales that Balzano's theory predicts may be particularly musically interesting. We do not consider here possible applications to atonal music, since the most commonly used configurations of the interface tend to relate harmonic structures to the diatonic scale and related scales (e.g. the pentatonic scale and harmonic minor scales), or to hypothesised microtonal analogues of the diatonic scale in the microtonal case.

The convention adopted for notating chord sequences is described in an appendix to this paper.

2 Balzano's theory

To better understand the interface, it is useful to outline informally the key result of Balzano's theory. For fuller accounts of the derivation, see Balzano (1980) or Holland (1989). The basic idea behind Balzano's theory can be summarised (for mathematicians) as follows: musicians with no special interest in mathematics may prefer to skip the next paragraph.

Balzano's theory identifies the 12-fold western pitch set, more or less irrespective of tuning, as being isomorphic with C_{12} , the cyclic group of order 12. This identification arises because the notes of the scale have a natural ordering (their pitch) which is circular under octave equivalence. The subgroups and group theoretic properties of C_{12} turn out to characterise many of the harmonic properties of the pitch-set very elegantly and economically.

Loosely speaking, from a more musical perspective, the key result is that most of the central relationships of tonal harmony can be expressed economically as simple, spatial relationships in an array (figure 1) in which notes of the chromatic scale are arranged as follows:

- 4 semitones (major thirds) on one axis,
- 3 semitones (minor thirds) on the other axis.

There is some evidence that this representation (the "thirds space" representation) is intimately connected with the way in which people perceive and process tonal harmony. Of course, there are other competing accounts; for example, see Watkins and Dyson (1985), Krumhansl (1991), Barucha and Krumhansl (1983), Rais (1992) and others.

Longuet-Higgins' theory was the principal initial inspiration for the interface, and many features have grown out of detailed consideration of aspects of Balzano's and Longuet-Higgins' theories, but no claim in this paper depends on the psychological accuracy or otherwise of either theory. However, it is *not* the case that any arbitrary spatial representation of the scale would serve our purposes equally well. Irrespective of the psychological worth of Balzano's theory, if we choose to characterise the pitch set as a set of 12 objects with circular ordering due to octave equivalence, then Balzano's result establishes that the third spaces provides the most economical co-ordinate space (in a sense explored below) for this pitch set. Let us now explore the ramifications of the result a little further.

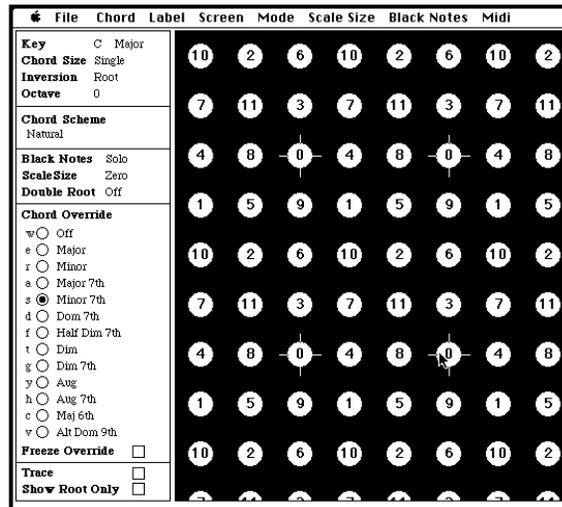


Figure 1 A very simple Harmony Space display. The labelling 0-11 refers to the 12 notes of the chromatic scale.

Balzano's theory may be viewed as resulting from a special case of a general strategy that is very common in the physical sciences. According to this strategy, for whatever objects are under consideration (notes, intervals etc), a complete, well-defined co-ordinate system is sought to span the objects. 'Complete' here means that all objects being considered must be specifiable using that co-ordinate system. 'Well-defined' means here that in any particular co-ordinate system, no object must be describable by more than one set of distinct co-ordinates. It is not hard to see that these are reasonable properties to demand of a co-ordinate system: for example, a candidate co-ordinate system for the euclidean plane would seem lacking in some way if it could not be used to specify certain points on the plane, or if several distinct co-ordinates were found equally well able to specify a single point.

Now, recall that since our approach characterises the tonal pitch set as a set of 12 objects with a circular ordering (under octave equivalence). This is not the only possible characterisation (since, for example, it ignores, frequency ratios) but it is a perfectly good one. From the viewpoint of seeking out complete, well-defined co-ordinate systems, Balzano's proof simply demonstrates that for such a pitch set, the following three structures:

- the chromatic ordering,
- the circle of fifths,
- and the two-dimensional thirds space,

correspond to three *complete* co-ordinate systems; and furthermore that these three co-ordinate systems are the *only* ones that provide *unique* co-ordinates for each note.

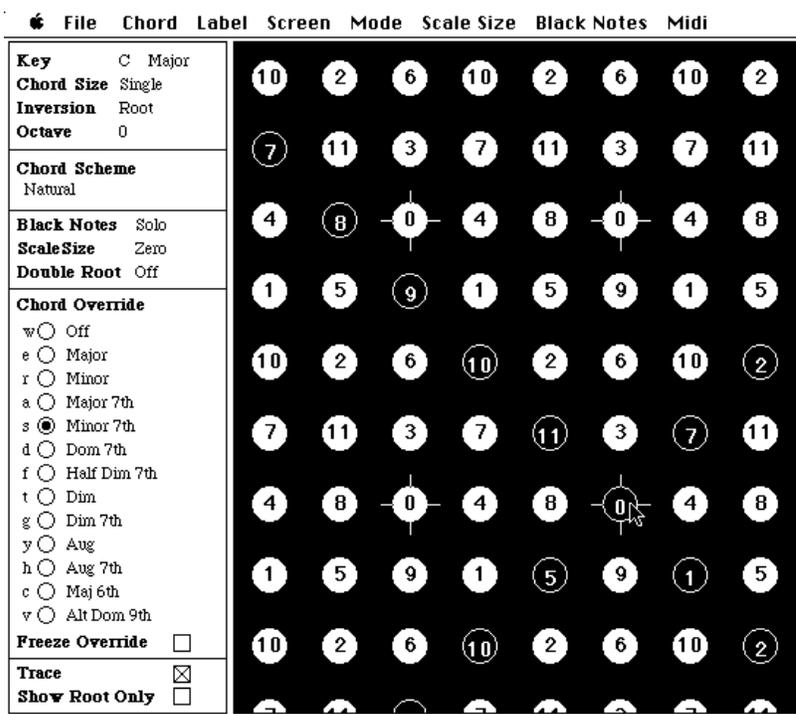


Figure 2. A Harmony Space display with notes labelled by semitone numbers. Two diagonal axes, which correspond to the chromatic scale and the circle of fifths are highlighted.

This result is quite strong, since it turns out that it applies not only to all possible 1- and 2-dimensional co-ordinate systems, but also to all possible n-dimensional systems for any integer n. Hence the result establishes that these three co-ordinate systems are especially privileged in the sense outlined above, compared with all other possible co-ordinate systems for describing relationships among a circularly ordered set of 12 notes. Let us now consider some of the characteristics that distinguish the three co-ordinate systems among each other.

Of the three co-ordinate systems identified by Balzano, the chromatic ordering seems particularly well-suited to representing melodic closeness (roughly speaking). The circle of fifths ordering is particularly well-suited to representing aspects of relationships between keys and certain common harmonic progressions. However, as we will see later on, the thirds space representation actually turns out to subsume the other two representations, since the chromatic scale and the circle of fifths appear as its diagonal axes (figure 2).

This is not to say that we are claiming that the thirds space is superior to all other representations of pitch for all purposes. This would be to neglect other factors that can be important. Indeed, other spatial representations may be viewed as combining two or more axes in ways especially suited to particular musical activities. For example, Common Music Notation may be viewed as combining, roughly speaking and amongst other things, a chromatic, melodic-related axis with a time axis. The guitar fretboard may be viewed as combining a chromatic, melodic-related axis with a fourths (or inverted fifths) axis, for reasons to do with harmonic relationships and constraints on finger dexterity. These 'representations' (i.e. music notation and the guitar fretboard) may be viewed as reflecting ingenious compromises between competing demands of harmony, melody, time, and physical, perceptual and cognitive limitations.

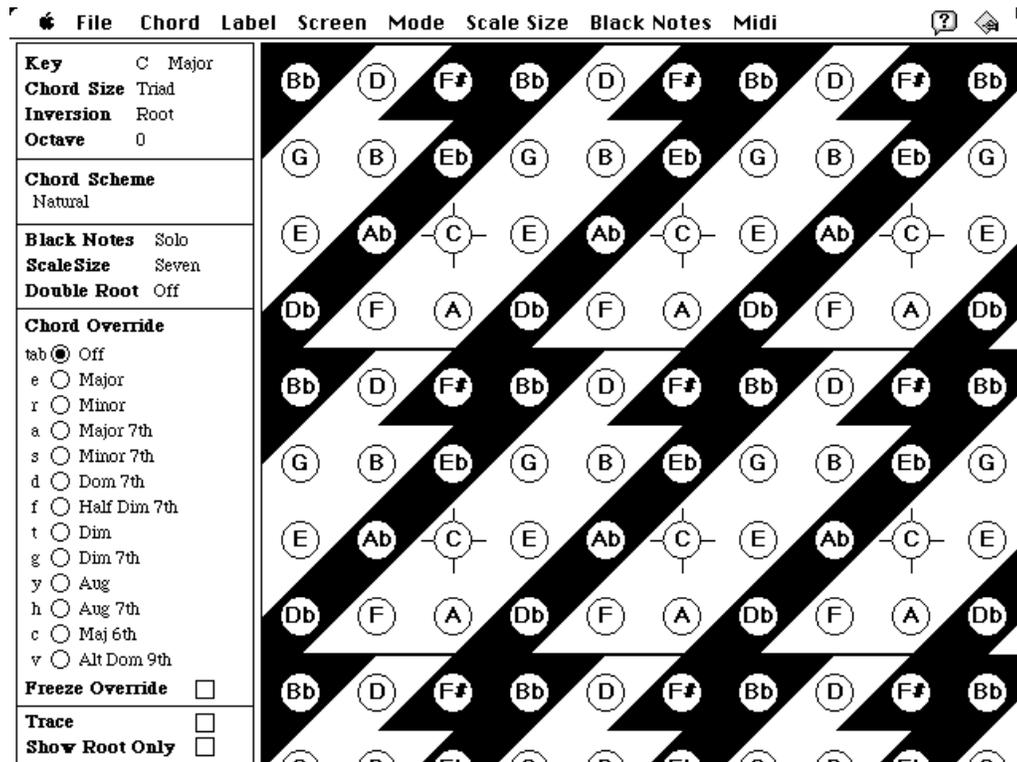


Figure 3. A Harmony Space display with the system in the key of C major.

The thirds space, on the other hand, may be viewed as using two dimensions solely to display harmonic relationships (i.e. simultaneous relationships between one or more notes). Thus it is not surprising that the thirds space appears to be uniquely powerful for exploring harmonic relationships. Other representations may be better suited to other purposes - e.g. analysing melodies, or exploring aspects of the interplay between melody and harmony. Of course, to encourage development of a rounded musical viewpoint, Harmony Space, like any other single representation, needs to be complemented with other representations.

The forgoing theoretical arguments can be illustrated at length with concrete examples. A detailed analysis of various musical phenomena and musical tasks seen from the point of view of a variety of spatial representation systems (including common music notation, pianoroll notation, the guitar fretboard and Longuet-Higgins' space) can be found in Holland (1989).

Of course, there are other possible ways of characterising the chromatic scale than as a set of 12 objects with a circular ordering: in some cases, other characterisations yield other complete, well-defined co-ordinate systems. For example Longuet-Higgins (1962) defines the pitch set in terms of overtone theory, yielding a different co-ordinate system that turn out to be intimately related to Balzano's. Different characterisations of pitch, and their co-ordinate systems tend to emphasise different aspects of tonal harmony. For an argument that Balzano's characterisation focuses particularly closely on aspects of immediate practical concern for novices, see Holland (1989).

3 Basics of the interface

We will now examine the interface itself. Let us start with the most basic version of the display from the current Macintosh version (figure 1). Notes of the chromatic scale are

numbered from 0-11 in ascending pitch order. Each circle represents a note. Clicking on the circle with the mouse pointer causes that circle to darken and the corresponding note to sound via an external MIDI synthesiser.

The horizontal and vertical axes of the grid correspond to movements in semitone step sizes 3 and 4, as motivated by the theoretical reasons noted above. The diagonal axes both turn out to have interesting properties. The NW axis yields a regular semitone ordering, as seen in figures 2 and 3. The NE axis corresponds to the circle of fifths, visible in figures 2 and 6. From an educational point of view, these are two important dimensions that western musicians constantly make use of (and novices need to learn about). They have simple physical, spatial interpretations in the interface. The NW axis is important since, for example, motion along it is much used in melodies, and relates to the notion of melodic 'closeness', as noted previously. The NE axis (i.e. the circle of fifths, based on regular semitone steps of size seven) is also very important in tonal music, for example being fundamental to the movement of commonplace bass lines, chord sequences, and key progressions. Alphabetical, semitone, roman numeral and various other forms of labelling are available, as illustrated in figures 2, 4 and 6. Due to limitations of screen space in the version shown, the note F#/Gb, for example, is always labelled F#. This is a superficial limitation easily overcome in versions with larger screen displays.

With the interface in its simplest state, all of the notes on the screen, when played as roots, normally sound in the same octave. This means that there are only really 12 distinguishable notes on the screen. The pattern of 12 notes has been propagated, like a wallpaper pattern, to allow certain other patterns, some of which we will examine in a moment, to strike the eye more easily. Notes can be shifted up or down one or more octaves interactively, using alternative mouse buttons (on the current SparcStation version) or alphanumeric keys (on the current version Macintosh version). To make changes in octave height visible as well as audible, the note circles are illuminated in different shades of gray depending on any octave displacement, as illustrated in figure 5. A data glove or high precision multiple-touch touchscreen might make a much better controller than a mouse, but these refinements must wait for later implementations.

3.1 The diatonic scale and other scales

By accessing the scale menu, it is possible to cause diagonal slices of any length from 1 to 12 to be picked out in white. This corresponds to picking out certain notes from the full set of 12 notes for special attention, i.e. forming scales. Diagonal slices on the NE axis are by no means the only way to form scales, but they are of particular importance from the point of view of Balzano's theory. Indeed, such scales have particular musical significance. For example, a diagonal section of length 5 taken from the circle of fifths corresponds to the pentatonic scale, while length 6 corresponds to the medieval scale, Guido's hexachord (Grove, 1940). By far the most important scale is the slice of length 7 the familiar diatonic scale, as shown in figures 3-7. With a scale of size 7, the diagonal slice appears visually folded alongside itself (for example as in figure 3), producing a repeating box shape with a chunk cut out. To make it easier for the eye to distinguish positions in the repeating pattern, the interface gives an option for the shading to be broken up by horizontal breaks, which can be seen in all of the figures except 5, 1 and 2. Notes in the white area correspond to the notes of the diatonic scale, and notes in the black area correspond to notes outside the diatonic scale. Note that in the special case of the key of C major, these areas correspond to the white notes and black notes on the piano respectively. We will refer to the pattern of white notes (the visible shape of the diatonic scale in the interface) as the 'key window'.

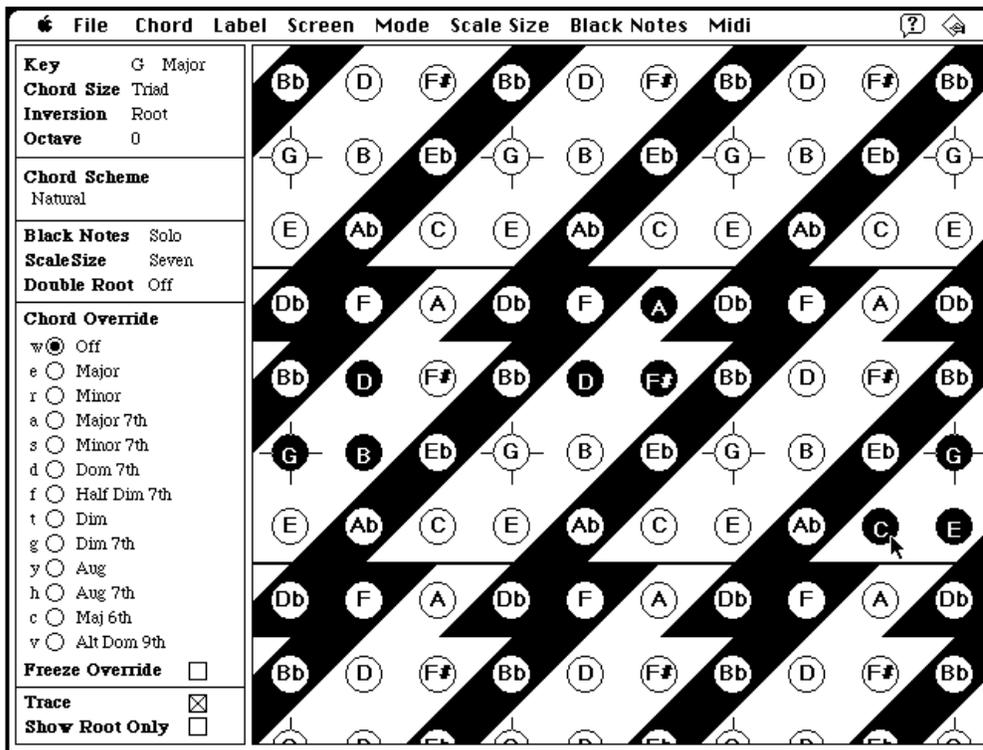


Figure 4. Three major triads in root position in the key of G. Normally, it would be convenient to play representatives of these triads closer together (e.g. all in a single key window). The triads have been triggered and displayed widely apart here for clarity of display on the printed page.

3.2 Key areas and modulation

Using the arrow keys on the alphanumeric keypad, we can move the key window (i.e. the white diatonic enclosure) around over the fixed grid of notes. This action corresponds to modulation. For example, if we slide the window up the diagonal cycle of fifths axis starting from the key of C, this corresponds to modulating to the dominant (figure 4). The key one 'notch' below our starting point on the diagonal axis is the subdominant (figure 5). There are 12 distinguishable positions for the key window: i.e. 12 major keys are distinguished. In general, keys that are musically closely related are nearby, and more remote keys are further away. We can use the interface as it stands to play simple melodies changing octave using mouse buttons or control keys as required. Note that if we run the mouse down the SE axis with the mouse button depressed, we will hear the chromatic scale (figure 2). If we run down the same axis but play only the notes in the key window, which is easily arranged by setting all the notes in the dark area to respond with silence, we will hear the diatonic scale (figure 3).

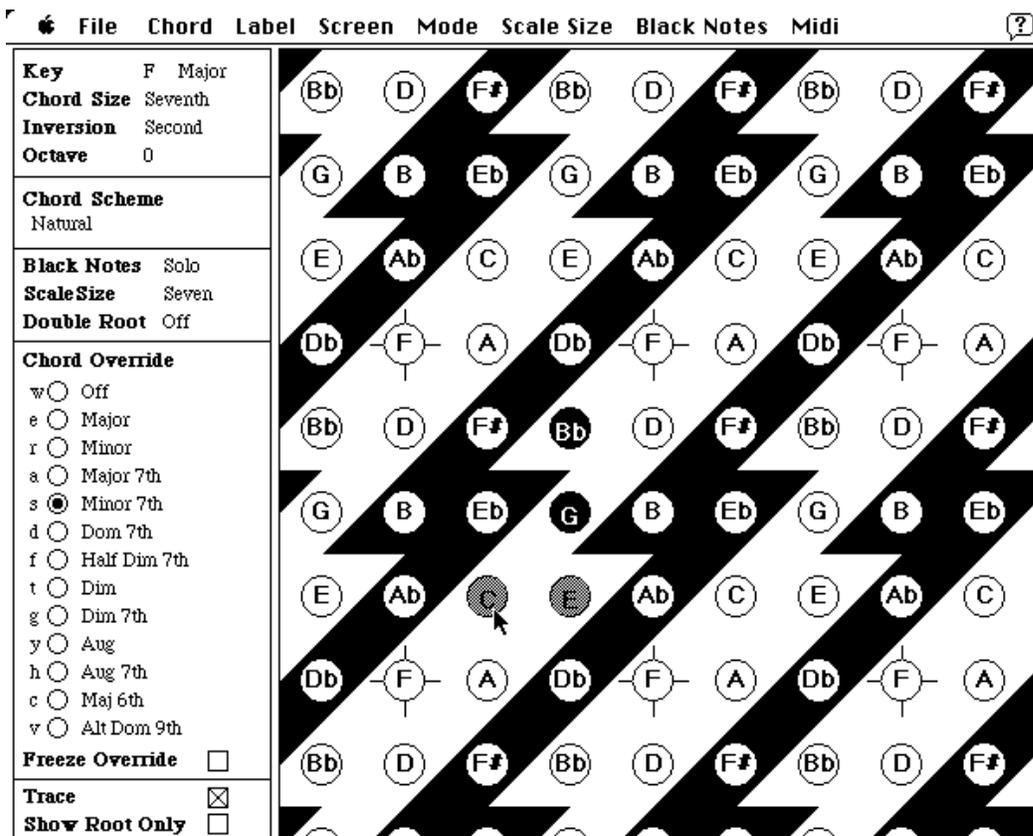


Figure 5. A C dominant 7th chord in second inversion in the key of F major. The fact that two notes at the bottom of the chord are grey (i.e. displaced up an octave) identifies that the chord is in second inversion. A slight variant of the display is illustrated (key windows are shown as continuous bands).

3.3 Chord construction

So far we have discussed only the playing of single notes. The interface allows the user to change the number of notes that will normally be played in response to a single mouse click, i.e. to play in triads, sevenths, ninths, etc, (figures 4 and 5) at any time by using a menu or control key. The interface can help beginners study the way that scale tone diads, triads, sevenths, ninths etc. build differently up forming different intervals on different degrees of the scale, independently of key transpositions. The rule for building chords on any degree of the scale is easy to appreciate visually. The rule is simply that the chord must expand along and up, and must fit into the key window. Intuitively speaking, the notes 'bounce upwards'. The visual behaviour gives the beginner a simple rationale for understanding why different degrees of the scale are associated with different heard chord qualities; namely, they conform to fit the shape of the key window: i.e. the diatonic scale.

3.4 Inversions, voicing and altered chords

In the simplest setting of the interface, all of the notes in a natural chord are sounded in root inversion in close position. In implementations with more than one mouse button the different mouse buttons allow root, first and second inversions to be played respectively (figure 5). Visually, in accordance with the metaphor for octave height already noted above,

notes played in different octaves from their root are shaded accordingly to show their displacement in the z-plane (figure 5). Amongst other things, this allows visual 'least move' considerations to be taken into account when moving between inversions.

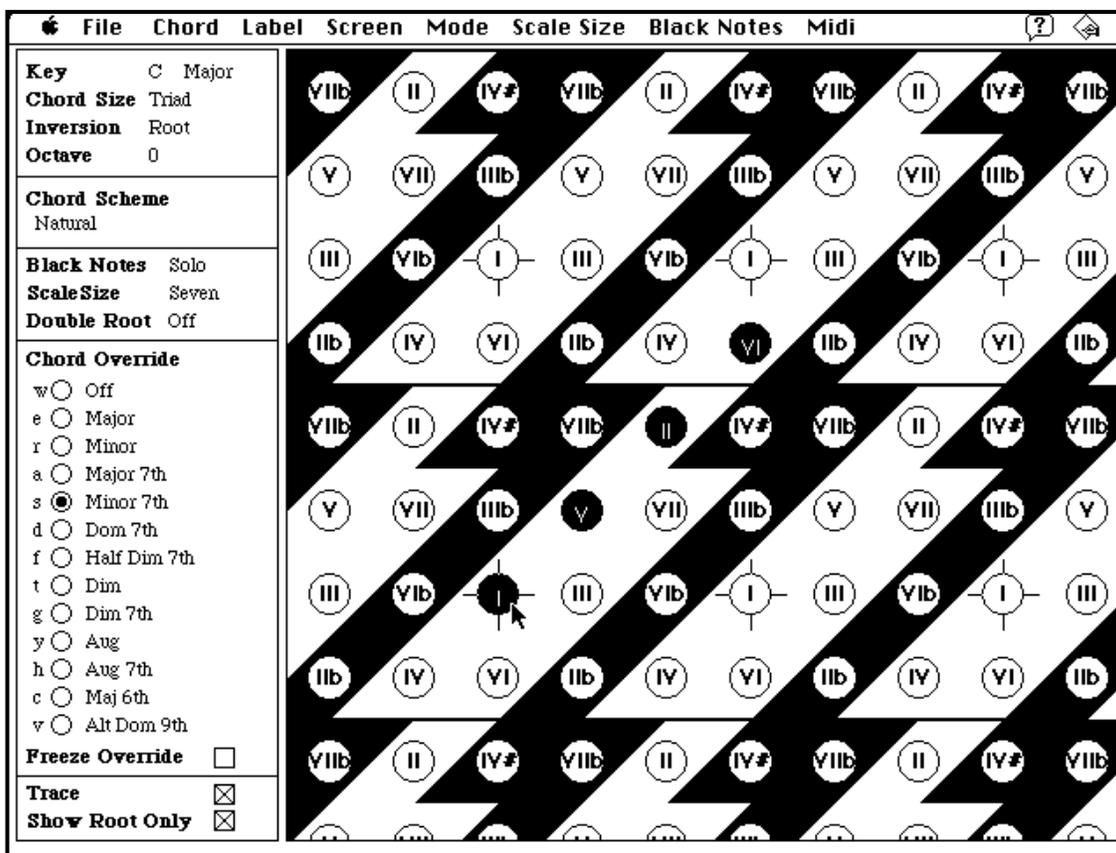


Figure 6 A Harmony Space display where notes are labelled with Roman numerals. The numerals move in step with the key window for modulations and changes of mode. The roots of a VI II VI progression are shown.

In addition to the moment-by-moment manual control of inversions, a chord editor can be used to pre-prepare families of chords voiced in any way desired: i.e. to displace particular notes by one or more octaves, double notes, suppress notes, etc. This allows appropriately voiced chords of whatever size (triad, seventh, etc) to be associated systematically with given degrees of the scale, or played individually on demand. So, for example, one can create and store a favourite Ravel or Debussy piano voicing scheme. Such schemes can be arranged to give good "least move" voice leading more or less automatically. Yet more complex voicing can be played by shifting dynamically between pre-prepared voicing schemes. More generally, the same chord editing mechanism that is used to set up appropriate voicings can also be used to associate arbitrary chords with any degree of the scale, for example, scales that involve altered chords, or minor or modal variants.

3.5 Triads and tonal centres

In Harmony Space, there is a one-to-one mapping between chord quality and visual chord shape, independently of key. Triads can be identified as the maximally compact three element shapes, i.e. shapes packed as closely together as possible, that will fit into the diatonic scale (figure 4). There can be seen to be exactly three chords of the first such shape (the major triads), and exactly three chords of the second such shape (the minor chords). There is only one chord of the third shape, (the diminished chord). Both major and minor chords are layed out in groups of three with a clearly visible central major chord and a clearly visible central minor chord. Note that this centrality is independent, in both the major and minor cases, of the way graphical boundaries are drawn. For example, in the minor case, there are just three 'naturally' occurring minor triads in the diatonic scale, arranged with roots a perfect fifth apart, one of which can be considered central, however the graphical boundaries are drawn.

The existence and location of the major and minor tonal centres are thus visually highlighted in a pedagogically useful way. This spatial metaphor for the two tonal centres reflects a hypothesis by Wilding-White (1961), as follows. Balzano (1980) quotes Wilding-White (1961) to the effect that the central notion behind a tonic is not a central note, but a central triad. Harmony Space can be used to explore and illustrate this hypothesis in a way which is easily grasped by beginners.

3.6 Accompaniment and harmonic analysis

Harmony Space can be used to analyse and play chord progressions. Indeed complete beginners can start making music straight away ignoring the theoretical discussions we have just entertained. Observations noted below, except where noted otherwise, refer to a formative evaluation with novices, reported in Holland (1989).

In a preliminary part of the evaluation, screening was used to select subjects with negligible previous musical training and negligible demonstrable musical ability: as a result of the screening, five subjects aged from 8-37 were chosen. Subjects had single sessions with an early prototype version of Harmony Space. Sessions lasted between 30 minutes and two and a half hours. Sessions were video-taped and audio-taped. The author acted as instructor. A range of tasks was taught individually including simple accompaniment and compositional tasks. Some of the informal findings of the formative evaluation were as follows.

The formative evaluation demonstrated that non-musically trained users (including a child) could use and make sense of the interface given initial guidance. Most of the users could be taught in a few minutes to play accompaniments to pieces in the major or minor mode in any key using typical two, three and four chord patterns. Some subjects were able to remember and find the location of tonal centres.

More extended 'classical harmonic' chord sequences could be played after a few minutes by most of the subjects, as diads, triads, sevenths or ninths. Such 'classical harmonic' chord sequences can be played by establishing a tonal centre, jumping to some non-tonic triad and then moving back to the tonic in a straight line through all intermediate triads along the cycle of fifths axis. This corresponds in Harmony Space to making single vertical straight line gestures (figure 6).

Some harmonically rich forms of tonal music typically use complex progressions that make use of moving targets (i.e modulations) in systematic ways. Harmony Space offers the opportunity of understanding such harmonic patterns strategically, in a series of trajectories, rather than just at chord-by-chord level. For example, the chord sequence for John Coltrane's 'Giant Steps', which modulates every two chords or so, is typically considered hard to play and to memorize. The Harmony Space trace makes it easy to visualise this piece as based on a very ingenious but essentially simple sequence of V I and II V I chords, i.e. straight lines,

modulating 'westwards' down in major thirds every two or three chords (figure 7). The sequence makes use of the "pun" that four such steps bring you back to the original key.

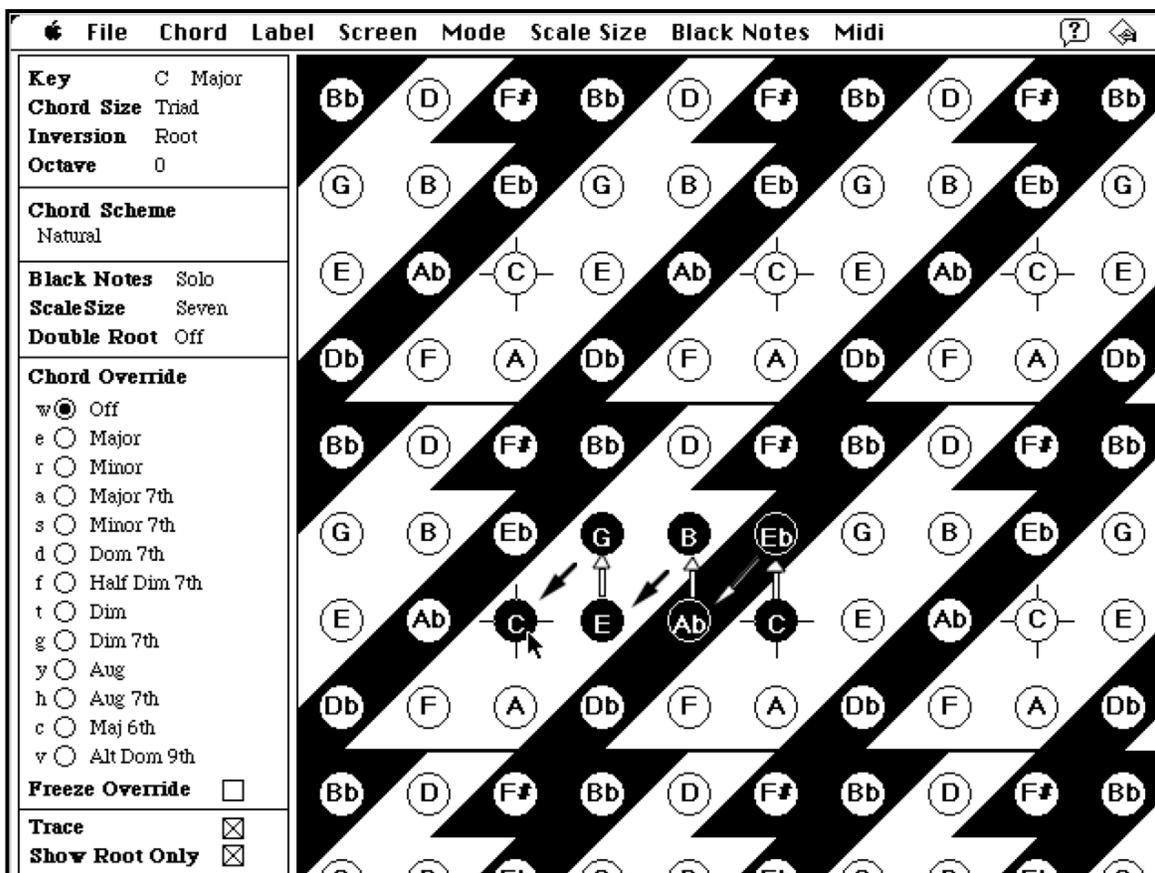


Figure 7. A simplified version of the sequence of chord roots on which 'Giant Steps' is based, as seen in Harmony Space (shown transposed into C major). The sequence modulates rapidly in such a way (not shown in this diagram) that it constitutes a series of V I chords, each in a new key. The full chord sequence has elaborations that include II V I sequences and modulations both up and down a major third. Each chord is a major or dominant ninth. Simplifications have been made to aid clarity on the printed page.

In a more recent and as yet unpublished pilot evaluation, for which data are currently being analysed, twelve year old volunteers from a comprehensive school in Milton Keynes received individual tuition from the author and others in Harmony Space after school. One interesting preliminary result was an informal demonstration that, given initial familiarisation with the interface, at least one child with minimal keyboard skills could be taught to play the chord sequences for 'Giant Steps' in Harmony Space in a minute or two.

Interestingly, in the case of musically educated adults using Harmony Space to analyse Giant Steps (Coates, 1992), once the trick used in this piece has been grasped, Harmony Space allows such users to explore what simple or radical variants on the basic idea might also work musically. For example, what happens if the V I's modulate consistently up a major third, or up

a minor third, or down a tone? Harmony Space makes it very easy to explore. Each idea presents problems to the ear, which then have to be solved, producing new potential pieces.

Contrasting uses of Harmony Space include ingenious harmonisation games devised recently by Dr Pat Howard of the Open University, as part of the second pilot evaluation (Howard, 1993). These games are playable by total novices to teach them to harmonise material such as Bach chorales.

Note that as currently designed, novices require guidance in the use of Harmony Space - there is no claim that simple exposure to Harmony Space without guidance would by itself teach novices.

4 Some educational uses of Harmony Space

Harmony Space has at least six general uses; musical instrument, particularly for sketching chord sequences; tool for musical analysis; tool for learning basic harmonisation; learning tool for exploring the theory of tonal music; discovery learning tool for composing and modifying chord sequences; and notation for aspects of chord sequences not obvious in conventional notations. Discussions of examples of each of these activities, as carried out by novices, can be found in Holland (1989).

5 Implementations of Harmony Space, and some variants

A family of variants of the interface have been implemented in prototype on the SparcStation. Each variant is specialised for various musical tasks. The basic implemented variants include: a microtonal version; a just intonation version; a version switchable to either Longuet-Higgins or Balzano configuration; a version for graphically performing harmonic analysis by graphical best fit; a version able to drive melodic and rhythmic figural patterns; a version as a recursive programming language able to drive the interface; a version designed to integrate with an intelligent tutor for music composition. Many other variants have been designed, prototyped or otherwise realised, including a human-powered performance-event version staged as a co-operative game at the Utrecht School of Art with Peter Desain and Henkjan Honing.

6 Related work

In many respects, Harmony Space is unique. Its relationship to various independently developed interfaces including Music Mouse, Harmony Grid, and Longuet-Higgins light organ is described in Holland (1989). A complete declarative model of the interface has been implemented as a logic program (Buchan 1991), which allows Harmony Space to be linked with a constraint-based musical planner (Holland, 1989) to form the core of an intelligent tutor for music composition.

7 Limitations

There are some aspects of harmony that the current version of Harmony Space does not represent well, for example, voice-leading and the visualisation and control of harmony in a metrical and rhythmic context. Harmony Space emphasises vertical aspects of harmony at the expense of horizontal aspects (using these terms as they are conventionally used in harmony - not with reference to the spatial layout of Harmony Space). Despite this, the Sparc version of Harmony Space can demonstrate several aspects of voice leading reasonably well using techniques sketched in Holland(1986). The current implementations of Harmony Space would benefit from various simple refinements including the following. On the current Macintosh version, though less so the Sparc version, melodic and rhythmic context is not provided by any

built-in features; rhythmic playing is hampered somewhat by response speed; only three or four levels of z-displacement are currently visible; the constituent notes of chords are not independently manipulable while sounding; mode changing, key changing, and label changing cause slight delays; inversions, voicings and non-standard chord scales should be more easily changeable. Other needed facilities include: a simple visual record playback facility, a MIDI-IN facility, more flexible ways of using the notes of a chord (e.g. via alberti patterns, arpeggiation patterns and finer manual control).

8 Further work

Around sixty further work projects related to Harmony Space are itemised in Holland (1989). We will note two or three of particular current interest. Harmony Space is based on theories that reconceptualise the domain in question, i.e. harmony. The reconceptualisations call for and give an opportunity to make radical changes to the curriculum and teaching approaches in this area. The developing of new curricula and teaching methods to take advantage of Harmony Space's theoretical basis has begun, but needs much more work. Various technical improvements are required. Various major variants of the interface await resources to implement, such as the full graphical analysis tool, and full associated discovery based composition tutors, currently only partly implemented.

9 Conclusions

AI and Education researchers have long been interested in taking AI theories (of particular domains, ways of learning, and styles of teaching) and of finding ways to exploit them for educational purposes. One less advertised way of doing this is to dispense with a teaching component, an expert component, or a student model, and to use something like the following strategy.

- Investigate as deeply as possible any cognitive theories or AI theories of the domain in question and gather computational models.
- Taking into account the best current HCI practice, try to devise an interface that represents abstract domain concepts in perceptually concrete, interactive, manipulable ways.
- Try to make concepts and relationships perceptually salient which should be conceptually important.
- Reduce cognitive load by representing explicitly and uniformly relationships that would otherwise have to be learned or calculated.

The ideas presented in this paper can be seen as example of this approach. We have outlined how cognitive theories of harmony can be used to design and build an interface that allows objects and processes of interest in harmony (e.g. notes, chords, chord progressions, key areas, modulations) to be manipulated directly using a single principled spatial metaphor to make a wide range of musical tasks accessible for novices. We have also sketched ways in which the interface makes certain expert tasks more tractable than using more orthodox tools and notations. We have shown how the two pointing devices together with visual and sound output are used to give rise to spatial, aural and kinaesthetic cues that all map uniformly into an underlying theory. Some recent implementations of Harmony Space have been discussed, together with some of the musical tasks which it facilitates.

Acknowledgements

This work has been assisted by grants from the University of Aberdeen Research Committee and the Open University Research Committee. Thanks to Mark Elsom-Cook, Tim

O'Shea and Mark Steedman for help early on. Thanks to Colin Watson, Fiona Houston, Matt Smith, James Bissett, Jim Buchan and Andrew Mayo for more recent work on various implementations. Thanks to Pat Howard, Daran Coates, Matt Smith and Nigel Morgan for recent useful discussions and work on evaluations. Thanks to Patte Wood, Heidi Kugler, Chris Chafe and all at CCRMA for having me for a brief but stimulating visit. Thanks in particular to Enrique Moreno for some arresting insights on Balzano's theory, and to Paul von Hippel for very helpful comments on this version.

References

- Balzano, G. J. (1980). The Group-theoretic Description of 12-fold and Microtonal Pitch systems. *Computer Music Journal* 4:4. Winter 1980.
- Barucha, J.J. and Krumhansl, C.L. (1983) The representation of harmonic structure in music: hierarchies of stability as a function of context. *Cognition* 13 63-102.
- Buchan, J. (1992) Initial steps in the design of an ITS for Music Composition. Unpublished MSc thesis, Department of Computing Science, University of Aberdeen, Scotland
- Coates, D. (1992) Personal communication about analysis and recomposition of 'Giant Steps' in Harmony Space, in the Department of Computing Science at the University of Aberdeen. Coates harmonically analysed Giant Steps using Harmony Space, and Holland performed various recompositions.
- Grove, G. (1940) (Ed. Colles, H.C.) *Grove's Dictionary of Music and Musicians*. MacMillan, London.
- Holland, S. (1992) "Interface Design for Empowerment: A case study from music" in Edwards, A. and Holland, S. (1992) (Eds) "Multimedia Interface Design in Education". Springer Verlag, Hiedelberg.
- Holland, S. and Elsom-Cook, M. (1990) "Architecture of a knowledge-based music tutor". In *Guided Discovery Tutoring*, Ed. Elsom-Cook, M. Paul Chapman Publishing Ltd, London.
- Holland, S. (1989) "Artificial Intelligence, Education and Music". PhD thesis published internally as OU IET CITE Report No. 88, July 1989.
- Holland, S. (1987) "Direct Manipulation tools for novices based on new cognitive theories of harmony". pp 182 -189 *Proceedings of 1987 International Computer Music Conference*.
- Holland, S. (1986). Design considerations for a human-computer interface using 12-tone three dimensional Harmony Space to aid novices to learn aspects of harmony and composition. CITE Report No. 7, Open University, Milton Keynes.
- Howard, P. (1993) Personal communication about harmonisation games devised by Dr Howard and played in a pilot evaluation at a local school in Milton Keynes in June 1993.
- Krumhansl, C.L. (1979) The psychological representation of musical pitch in a tonal context. *Cognitive Psychology*, 1979, 11, 346-374.
- Krumhansl, C.L. (1991) Memory for musical surface. *Memory and Cognition* 19,4, pp 401-411.
- Longuet-Higgins, H.C., (1962a); Letter to a Musical Friend. *Music Review*, August 1962 pp. 244-248
- Longuet-Higgins, H.C., (1962b) Second letter to a Musical Friend. *Music Review*, Nov. 1962.
- Mehegan, J. (1959) *Jazz Improvisation I: Tonal and Rhythmic Principles*. Watson-Guption Publications, New York.
- Rais, M. (1992) Jaan Soonvald and his musical system. *Leonardo Music Journal*, Vol 2, No 1 pp 45-47.

- Watkins, J. and Dyson, Mary C. (1985) On the perception of tone sequences and melodies. In Howell, Peter, Cross, Ian, and West, Robert (Eds.) *Musical Structure and Cognition*. Academic Press, London.
- Wilding-White, R. (1961) "Tonality and Scale theory". *Journal of Music Theory* 5:275-286.

Appendix

The convention used to notate chords is that of Mehegan (1959), where full details of the notation can be found. The key points are summarised below. This notation is particularly well suited to describing and reasoning about chord sequences in the context of Harmony Space. In the Mehegan system, chords are described using roman numerals, though this does not imply that the rules of classical functional harmony necessarily apply. A sequence of numerals is taken to consist of triads, unless the texture of the sequence as a whole is specified to be otherwise, and except for any individual chords annotated differently. A sequence of numerals as a whole may be declared to refer to sevenths or ninths or diads (again, excepting any individually annotated chords).

Upper case roman numerals are used throughout the system: the case (i.e. upper case or lower case) of a roman numeral does not imply anything about chord quality. Unless a roman numeral is explicitly annotated to indicate additions, alterations or a specific quality, the corresponding chord is understood to be of 'scale-tone chord quality', a term which is explained below. Chords of scale tone quality are constructed by stacking thirds upwards from the root so as to stay inside the prevailing key or mode. Note that in the case of major and harmonic minor scales, and in modes that share the same notes as these scales, this is a well-defined method of chord construction. Annotations for inversions, altered qualities, etc exist in this system but are not described here as they are not used in this paper.