

# An Investigation into the Tuition of Music Theory using Empirical Modelling

0503985

## Abstract

Music theory is a subject that is often thought essential to the learning of a musical instrument, but is fraught with complexities, contradictions and subjective thought. The human ear is often not able to accurately discern pitch at a level necessary for a fluent understanding of harmony at the beginning of tuition, making it harder to gain understanding of the concepts involved. This paper explores using Empirical Modelling to create an open-ended visualisation environment for the tuition of basic music theory, in order to aid general musical understanding.

## 1 Introduction

Music theory is a difficult subject - no musician would deny that. It can take years of practice and study to gain an understanding of the concepts involved and to be able to apply that knowledge to an instrument or in composition.

Empirical Modelling (EM) is a field of study with a great affinity for application in Educational Technology. Its focus on “thinking with the computer” allows for strong constructionist learning, an essential tool for music. By applying EM techniques to music theory visualisation, this paper hopes to explore how the computer can aid learning in this context. This paper will first discuss the relevant research and background for music theory, educational technology and Empirical Modelling. The subsequent sections will describe the model created, and an analysis of the model and how well it meets its goals.

## 2 Research

### 2.1 Music

Music is one of the oldest and oddest of humanities creations. At its heart, music has a number of core philosophical questions - for example, why do we find certain combinations of sounds pleasing, what is it about these sounds that evokes dif-

ferent emotions? Despite the lack of definitive answers, a vast body of knowledge has evolved to describe the interactions of notes and combinations of notes in an attempt to create a set of rules which can be used to both analyse and create. Despite being built on weak philosophical foundations (for music theory is arguably an attempt to make axiomatic that which is both subjective and ill-defined) this knowledge has taken music to ever higher levels of complexity.

Music theory can be an incredibly involved and deep subject - many musicians can talk in a language which makes mathematical jargon seem simple. However, it is at the basic level that its effects can most be seen. The leap in compositional ability of trainee musicians upon learning chords, scales, and intervals is indisputable. While music theory must be taught with appreciation of one condition - all rules are subjective and made to be broken - it can still aid in an appreciation of composition.

Since it would be almost impossible to encode all the possible permutations of music theory into a visualisation environment let alone find a way to navigate through the subjectivities inherent in such a field, a subset of basic theory must be decided on before design can commence.

A good working definition of music for this paper is given by Lerdahl and Jackendoff (1):

We take the goal of a theory of music to be a formal description of the musical intuitions of an educated listener.

They then go on to define musical intuitions as

the largely unconscious knowledge which a listener brings to music and which allows him to organize musical sounds into coherent patterns.

This idea of organisation sheds some light on the problem of finding a basic subset - which parts of music theory describe the basis of these organisations, that is, which parts can be applied to almost all music. We immediately hit upon the problem that there is no such set, as music has evolved into some very varied forms. However, the most common uses 12 tones to an octave, and is found in both Chinese and Western music. Indeed, Western music has been shown to create recognisable emotional responses universally (2). Therefore this is the basis used for the rest of the paper. Using this, there is a clear set of basic music theory that must be learnt before all else which can be used to describe almost all music. This is the definition of scales, chords and intervals.

The notes are labelled A to G, with a number of accidentals, which are semitone shifts, giving 12 notes: A, A#, B, C, C#, D, D#, E, F, F#, G, G#. Note that many of these can be written in different ways, for example A# is the same note as Bb. This paper will write all notes as given in the list above for clarity's sake.

Scales are defined as a series of intervals. The Major scale, for example, first rises 2 semitones, then 2 again, and so on, giving the pattern of semitone shifts of 2-2-1-2-2-2-1. By picking any root note, the scale can be constructed using these intervals. Picking C gives the C-Major scale of C-D-E-F-G-A-B. Scales are a useful tool as most pieces of music can be conveniently defined by scales.

Using scales, it is then possible to define chords, or sets of notes to be played together. These can be defined in terms of scale degree; taking the first, third and fifth notes of a major scale gives a major chord (for example, C-E-G is

C-Major). A chord can also be defined in terms of intervals. A major chord is the root note, a note one major third up (4 semitones), and a note one perfect fifth up (7 semitones).

Having defined these simple rules, a basic music theory begins to take shape. By choosing a key (eg, C Major), a composer can define a set of notes which he knows will fit both the "feel" of that key (the major key is often described as "happy", though as with all music theory, this does not apply pervasively) and which will work with each other without sounding atonal. Any chords selected using these notes will fit in the feel of the song and work with a melody selected from these notes.

We already therefore have a strong toolbox for describing the organisation of a piece of music. The complexities of music theory begin to enter when we consider the interplay of notes within the scale, different scale types, and different chord types. It is at this point that constructionist learning can become very important in understanding music theory.

## 2.2 Educational Technology and Constructionist Learning

While the term "Educational Technology" can apply to a wide variety of technologies, this essay will discuss it purely in the context of learning using computers. Educational Technology is a field that has come into its own with the advent of personal computers. With high levels of availability and basic technical knowledge, computers are able to impact education positively in many ways.

One of the most interesting uses of computers in this field is with respect to constructionist learning - the idea of "learning by doing." By actively participating in a field on their own terms, rather than being taught its principles directly, students construct their own knowledge of the field based on experience. Constructionist learning is becoming more and more accepted and promotes a deeper, more thorough understanding (3). Constructionism limits the role of a teacher to that of a guide - imparting information in such

a way as to promote individual exploration of a subject, and to push students towards certain realisations and achievements.

Examples of this abound. We don't learn to ride a bike by being told by a teacher how to balance, and piano tuition is done while sitting in front of a piano, actually performing and practicing. While there are obvious differences between these and music theory (they contain an element of physical manifestation, in that the body has to be taught how to move), the implications are clear. Promoting constructionist learning of music theory is difficult however as for a beginner the ear is often not developed enough to be able to discern pitches accurately. One of the most challenging things for a beginner guitarist is learning to tune the instrument! It is clearly hard for a beginner to explore how notes fit together when this ability to appreciate these relations is diminished.

Clearly any tool to aid in this understanding cannot be all encompassing; merely visualising the relationships discards the element of music which makes it most appealing, that of the subjective auditory response. Any tool to visualise the context of music theory can therefore only act as a guide, an environment in which certain rules are set but which allows the user to explore the boundaries of these rules.

### 2.3 Empirical Modelling

Empirical Modelling as a field of research offers clear benefits in this situation. Empirical Modelling promotes the use of "thinking with computers," using the computer to aid thought development and explore a modelled universe. It is easy to see the parallels with constructionist learning - by creating an EM model of basic music theory, a user can then explore how these concepts interact by redefining various observables and building on the results. Moreover, since the theory encoded in the EM model is directly configurable by the user, a trainee musician can bend the rules to his whim and see how it affects the visualised tonality.

Empirical Modelling encourages the use of

programming models through dependencies and observables. Music theory itself is easily interpreted as a set of dependencies - at its simplest level, it describes how the tonality of various notes depends on the context of the piece, and gives a model for describing music based on this set of interacting dependencies.

Finally, Empirical Modelling attempts to bridge the subjective gap between user and computer. It explores the notion that each user of a program or model has their own, unique subjective experience of their usage, and traditional computer programming and modelling does not take this into account. Music theory, while attempting a semi-objective descriptive language for music, still has many subjective elements - including, for example, the tonality that the model in this paper attempts to visualise. By using Empirical Modelling over conventional methods, it is hoped that this subjective divide is bridged.

### 2.4 Visualisation

There has been varied research on the issue of how to visualise musical concepts. (4) considers assigning colour based on frequency of the note (based on work in (5)). Indeed, this exploits the similarity of colour and musical tones; both are entirely definable by wavelength. However, this is rejected (since the one-to-one mapping is fairly bereft of information content) in favour of an author-defined, arbitrary colour scheme which the author admits is based on emotional response. (6) takes another approach; visualising tones on a grid, whereby distance between notes implies the level of relative consonance. Their approach however is slightly different - attempting to visualise recordings of music temporally as opposed to merely viewing how musically theoretical constraints affect tonality. (7) provides an interesting study on the association of colour with emotion, as triggered by music, concluding that certain hues and colours are strongly associated with certain musical "feelings." Many composers suffering from synaesthesia would strongly associate colours with mu-

sical keys (8).

There are clearly many ways to visualise music, and so the driving force on how to decide on a mechanism is to try accurately represent to those concepts important in basic tonality. There are three overall aspects to consider: the emotion or feel of a set of tones, the distance of tones or sets of tones from each other, and the relative tonality (or atonality) of sets of tones.

The emotion, or feel, of a set of tones is explored in (4) and (7). The visualisation tool described here must be able to represent this feel, and how it is affected by the tones selected within a key. For example, even though C Major may be visualised using a happy, bright colour, the minor chords contained within, if played on their own, would be visualised as relatively sombre, sad sounding. This element of the visualisation is the most subjective, and therefore the most challenging to represent.

The distance of tones or sets of tones from each other is an important aspect of music visualisation. While almost everyone can appreciate that different notes have different distances between them, it takes rigorous training to be able to recognise a major third (for example). Visualising the distance between notes in chords and scales can help give a strong understanding of theory.

Finally, the relative tonality of notes is extremely important. Two notes played a single tone apart can sound very dissonant, but when used inside a 7th chord add only a small amount of dissonance to create a jazzy “tang.” Indeed, 7th chords form the basis of much of Jazz music. Thirds on the other hand are extremely consonant, and indeed major and minor chords are entirely made of thirds (a major chord is a minor third stacked on a major third). By being able to visualise the relative dissonance of intervals the student can start to learn how to construct their own chords or use different intervals to create different feels.

### 3 Creating The Model

The model created contains three main elements in its display - a piano roll, a scale column and a chord column. The piano roll is mainly to add context to the other elements; it shows and labels two octaves of a piano and marks the notes from the currently selected key. The key plays a minor role within the model - it is merely to place the scales and chord visualised into the proper context. For example, visualising C Dorian in the key of D Major to see how the component notes and tonalities differ, despite both being scales starting with D (see figure 1).

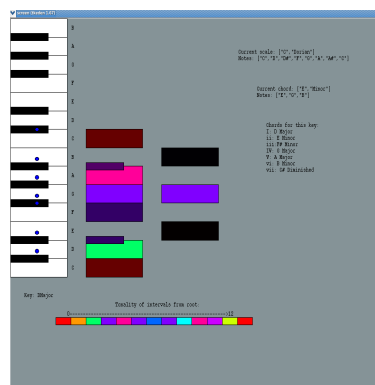


Figure 1: Figure showing how C Dorian is visualised in the key of D Major. Also shown is the chord of E Minor. Note that although E Minor is within the key, the notes that are not in the current scale are darkened.

The scale column displays the notes of the currently selected scale, coloured with their calculated tonality. The chord column is identical, except for displaying the currently selected chord. By changing the scales and chords, the user can see how the notes fit together and (most importantly for the beginner!) which notes are common to which constructions.

There was one major issue during development of the model. The decision was made to visualise each of the three axes described above as a hue, saturation, and intensity triple. The reasoning behind this was that these three axes are easier to separate out visually, than, say, attempting to see how much green or red is in a colour.

However, after implementing this it became

clear that there was simply too much information to fit in; by setting saturation to the tonality of a note relative to its nearest neighbours in the scale or chord, the colours become too washed out and similar to visualise the tonality of the note with the root (which was represented by its hue). This is shown in figure 2:




Figure 2: Washed out colours on the scale representation.

To combat this, the decision was taken to visualise only the tonality of the note from the root of the chord or scale, and whether it is in the current scale (for the chord) or key (for the scale). While this limits the functionality of the model, these intervals are more important for visualising a construction since it is these that create the feel of a chord. This modified visualisation is shown in figure 3:




Figure 3: Final, clear visualisation

The hue of these intervals is based on a tonality array within the model, containing numbers

in the range [0,1], 0 being entirely atonal and 1 being tonal. These have been set initially by the author and in this a sense of subjectivity creeps in. However, the values set create a set of distinct and bright hues that are easily identifiable once the association has been made with that colour, thus allowing the model to conform to the original goals. Moreover, true to the goals of EM, the user can redefine these tonalities to whatever value he likes. This might take the form of, for example, attempting to visualise some of the tonalities of middle eastern music, in which the set of most commonly used intervals is very different. A comparison of how the tonality array affects the visualisation is shown in figure 4:

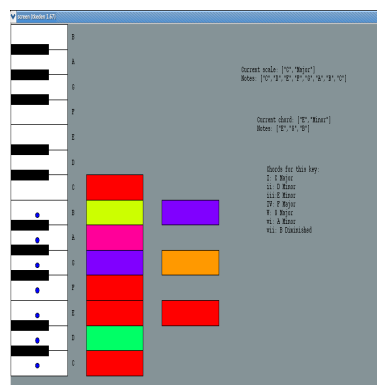


Figure 4: Scale of C Major visualised using arbitrarily chosen tonality values

One of the key elements of music is its temporal evolution - a static chord within a scale makes for a very uninteresting composition. Therefore any tool for visualising theory needs to build in an appreciation of how sequences of chords or scales relate to each other and the key. To this end, the model has a built in animation framework, with which the user can add a sequence of chords and scales, to be displayed for a specified length of time. This allows, for example, a user to input the chords to their favourite song, and see clearly how the chords fit with each other. Finally, the model shows the chords that can be generated from the current scale using roman numeral notation. This allows the user to quickly see the chords that are likely to be most tonal given the scale, and then explore these chords

and why they work.

A small subset of the most common scales and chords are included in the model. The user can easily add more, based on any list of intervals. In this way the user can create custom chords and scales and attempt to visualise how they interact. An interesting side-note to this is that since the list of chords generated from the scale is dynamic it can display some very odd chords that work with the current scale. Indeed, in testing the model the author has already found some chord/scale combinations that don't seem like they should work together!

To keep the display synchronised with the backend music theory model, there are a set of dependencies defined for the current key, scale, and chord variables. When the user manipulates these, the new tonalities are calculated and the display updated. The animated element relies on these dependencies - it sets up a timer which changes the current chord or scale to the correct value at each time step. However, beyond these the interaction is text based. For example, to view the list of defined chords or scales the user would query the respective variables, and then can change the current scale or chord based on these.

## 4 Analysis and Conclusions

### 4.1 Utility of model

The model created sets up a basic framework for the manipulation of chords and scales relative to each other. In this sense, it performs extremely well, allowing the user to manipulate the displayed musical context and visualise the relationships of the component notes. During testing, the author found that the relative placement of the notes helped inform understanding as much as the tonality-based colour relationships. In doing this, the model has shown that it fulfills the brief of "thinking with computers" extremely well - the model acts as an intuitive aid for the understanding of music theory.

The brain is obviously very adept at spacial reasoning, and so being able to place abstract au-

dio information such as pitch in a 2-dimensional visual context brings the full force of the brain's spacial power to bear on the task at hand.

The colour visualisation is slightly less intuitive than hoped, but with a small amount of time and commitment the relationships between the displayed colour and the tonality becomes clarified. The flexibility of the model is its strongest feature. All elements of the basic music theory encoded can be manipulated, allowing the user to explore a universe of abstract relationships - which can be entirely alien to the established musical relationships of Western music if the user so wishes! By setting the visualised tonality to the user-definable tonality array, the subjectivity of the author is removed and instead the trainee musician is invited to play with and bend the established rules.

### 4.2 Limitations and extensions

There are a number of potential drawbacks with the model as currently developed. Firstly, it lacks a goal or guiding philosophy. It simply presents the environment as is, and lets the user redefine the rules and applications of the rules so that they can see how that action affects the interactions of these constraints. However, there is no process that tells the user how to get more out of it; there is no evaluation or analysis given to the user's exploration. The user may easily learn the relationships between the notes of, say, C Major and the A Minor scale, and see that one is the relative minor of the other, but what that means musically is perhaps ambiguous. The separation of perception inherent in experiencing something visually and experiencing it through sound is still somewhat apparent in the model. The hope however is that the student gains an appreciation not of what the rules sound like, but why the rules are so defined - for example, it is abundantly clear, when using the model, why the chord of C Major sounds atonal when played in the scale of D Major - they have only one common note, and that is displayed as only slightly consonant! When the user then returns to the instrument, the lessons learnt in this exploration can then be applied.

However this discussion presents the obvious possibility that extending the model to include sound synthesis of the visualised concepts would be extremely beneficial. Doing so would very quickly associated the visualised tonality with how the user actually perceives these intervals, and would likely make interpretation (and thus learning) much faster.

A number of other extensions also suggest themselves. The model merely presents information and allows manipulation of it. Giving the model a guiding philosophy, perhaps through a set of exercises for the user or through an eventual goal (or test) would give the model more clearly defined purpose. For example, setting the user the task of working out which modes have the same notes over the whole octave, and then working out which chord types are most tonal within each mode starting with the same note, would give the user an extremely strong appreciation of how modes work - usually a difficult subject to get ones head around!

A final, more obvious extension to the model is to simply extend the corpus of music theory encoded within it - perhaps by introducing some rules from harmonic theory, or by extending the temporal function to allow melodies and complex rhythms. However, such an extension would run the risk of diluting the nature of the model - its simplicity is a strength, not muddying the novice's mind with irrelevant (at this stage of learning) facts that would run the risk of confusing rather than informing.

### 4.3 Conclusions

This paper and associated modelling study have explored the idea of a visualisation environment for exploring music theory through constructionist learning. Music theory is an obvious candidate for this form of learning, since a major component of tuition on instruments is constructionist. Since our sense of sight is strongest amongst all of them, our brain is extremely adept at reasoning spatially and visually. Ears, on the other hand, generally need rigorous training to give them the ability to accurately discern individ-

ual intervals. This can make it harder to learn music theory at a beginner level, as the brain is unable to clearly hear the relationships between the individual rules set out. By visualising these using Empirical Modelling, it is clear that beginner musicians can be effectively aided - indeed, the author came across many interesting and previously mis-understood musical relationships during development of the associated modelling study.

While the study described here functions more as a proof-of-concept than a fully fledged tool for educational technology, it shows the potential power of Empirical Modelling in aiding understanding of advanced concepts. The initial hope of displaying three separate axes of tonality (intra-note tonality, note-to-root tonality and whether the note is within the scale/key context) proved over-ambitious, but by removing one of these axes the model still provides a strong colour association for tonality, and the unforeseen advantages of visualising the notes spatially helps overcome this limitation. Overall, this model powerfully demonstrates the potential of Empirical Modelling as a tool for aiding musical understanding.

## References

- [1] F. Lerdahl and R. Jackendoff, *A Generative Theory of Tonal Music*. MIT Press, Cambridge, Massachusetts, 1983.
- [2] T. Fritz, S. Jentschke, N. Gosselin, D. Sammler, I. Peretz, R. Turner, A. Friederici, and S. Koelsch, "Universal recognition of three basic emotions in music," *Current Biology*, 2009.
- [3] J. Brooks and M. Brooks, *In Search of Understanding. The Case for Constructivist Classrooms*. Association for Supervision and Curriculum Development, 1999.
- [4] Unknown, "Synesthesia for all crafting a colour model of schubert's 'erlknig'," 2007.

- [5] K. Giannakis and M. Smith, “Auditory-visual associations for music compositional processes: A survey,” 2000.
- [6] T. Bergstrom, K. Karahalios, and J. Hart, “Isochords: Visualizing structure in music,” *Proceedings of Graphical Interfaces*, 2007.
- [7] R. Bresin, “What is the colour of that music performance?,”
- [8] R. Cytowic, *Synesthesia: A Union of the Senses*. MIT Press, Cambridge, Massachusetts, 2002.