Hawkett, Andrew

An empirical investigation of the concept of memes in music using mass data analysis of string quartets

AN EMPIRICAL INVESTIGATION OF THE CONCEPT OF MEMES IN MUSIC USING MASS DATA ANALYSIS OF STRING QUARTETS

ANDREW HAWKETT

A thesis submitted to the University of Huddersfield in partial fulfilment of the requirements for the degree of Doctor of Philosophy

The University of Huddersfield

March 2013

Abstract

Dawkins introduced the concept of the meme as the cultural equivalent to the gene (1989, pp. 189-201). To illustrate the concept, Dawkins cited 'tunes, ideas, catch-phrases, clothes, fashions, ways of making pots or of building arches' (1989, p. 192) as examples of memes. All of Dawkins' examples are elements of culture that have evolved over time. Because music is a part of culture, then under Dawkins' hypothesis, memes should exist in music. After all, the first of Dawkins' examples was a 'tune'.

Jan expanded on Dawkins' ideas with a thorough investigation into memes in music (2007). This was done on a number of different levels within music, from melodic lines to overall structure, using a range of examples within music. Whilst providing a strong case for memes, Jan was not able to provide evidence from an analysis encompassing a large dataset of music. However, Jan does provide a number of possible methodologies for analysing memes in music, including investigating memes across time periods using single lines of notes (2007, p. 211). The present research expands on Jan's suggested methodology by looking at short monophonic three- to eleven-note patterns in music across five different non-traditional musicological time periods within a large dataset of string quartets.

A search for memes in music is conducted using a range of scores. These are converted to MusicXML documents, which are then imported into a relational database. Data mining is then implemented on the resultant dataset to produce a series of ranking positions for monophonic note patterns within the music based upon the relative frequencies of their appearances within specified time periods. Additionally, a similarity algorithm is used to investigate the possible ancestral relationships between different monophonic note patterns. Within the limitations of the working definitions and assumptions made in the research, it was shown that there is evidence for the evolutionary properties of selection, replication and variation, and the replicator properties of longevity, fecundity and copying fidelity for some monophonic note patterns within the dataset.

Contents

Word Count: 84,316

List of Figures

List of Tables

List of Examples

Acknowledgements

I am deeply indebted to my supervisor, Steven Jan, for his support, encouragement and especially for his depth of knowledge and understanding of both music and memes. I am also deeply grateful for the patience of my partner, Martyn Prigmore, and for his support and advice, especially in the mathematical and technical sections of this thesis.

Abbreviations and Formatting

Abbreviations

The following non-standard abbreviations are used within this thesis:

A-PT	Absolute pitch value excluding all durations pattern type
AD-PT	Absolute pitch value including relative durations pattern type
PTD	Pattern type descriptor
PTI	Pattern type instance
PTL	Pattern type length
R-PT	Relative intervallic pitch excluding all durations pattern type
RD-PT	Relative intervallic pitch including relative durations pattern type

A full explanation of these abbreviations is given at the end of the introduction to Chapter 5.

Formatting

The commercial product names referred to in this thesis are given in *italics* whereas non-commercial product names are given in normal font. The table and column names used in the database are given in the Arial font.

The first statement of a working definition used within the main text is highlighted in ***Blue***. Each working definition is also listed in Appendix 1.

Introduction

Richard Dawkins argues that evolutionary forces drive cultural change (1989, p. 192). In, *The Selfish Gene* (1989) he argues that evolution is not restricted to biology but can exist in other areas as well. From this, he concludes that culture can evolve. In order to apply evolution to culture, he introduces the concept of the 'meme' that he believes possesses, in cultural terms, roughly equivalent properties as the gene (1989, pp. 189-201). As examples of memes he gives 'tunes, ideas, catch-phrases, clothes fashions, ways of making pots or of building arches' (1989, p. 192). Memes pass themselves between brains and, as a result, spread through society (Dawkins, 1989, p. 192). If it is accepted that music is a part of culture, then following Dawkins' ideas on memes being a cultural equivalent of a gene, music must also be made up of memes and these memes must be subject to evolutionary forces.

Listening to music from different periods in time there can be no doubt that changes have taken place. For example, a brief comparison between Gregorian Chant and Rutter's *Te Deum* reveals that differences include the range of instrumental and vocal parts involved, overall structures, harmonic palates, melodic structure, rhythmic complexities etc. However, are these changes attributable to evolutionary processes acting on memes, i.e., do memes exist in music?

Jan (2007) provided a strong argument for the existence of memes in music by hypothesising that memes could exist at a number of different levels within music ranging from short motifs to structural patterns. A number of examples within music were investigated by Jan to help illustrate the different forms that memes could take within the different levels. Following on from Jan's work, an attempt is made here to provide further evidence to support the meme hypothesis in music by looking for empirical evidence for memes using mass data analysis. The investigation uses as a basis one of Jan's suggested methodologies for analysing memes in music at the inter-work level, in that it uses 'horizontal … segmentation to generate a candidate pattern' (2007, p. 211) that allows the 'investigation of a single meme and its transmission-mutation history within a dialect over space and time' (2007, p. 211), albeit here with many candidate patterns simultaneously (although no attempt is made to use statistical modelling to determine overall numbers of candidate memes within the corpus, or music in general).

There are, however, various problems with investigating memes in music. Part I argues that there is no general consensus on what constitutes a meme, which is partly a reflection of there being no general consensus on what constitutes culture. What is agreed by all theorists is that memes, if they exist, evolve, and it is hypothesised in the present research memes evolve through natural selection. A brief investigation of what constitutes evolution by natural selection follows, and concludes that three processes are required: selection, replication and variation. Additionally, Dawkins' replicator properties of longevity, fecundity and copying-fidelity are also explored (1989, pp. 193-194). Consequently, a definition of a meme is alighted upon which is broad enough to encompass memes in music, whilst including references to the evolutionary processes and replicator properties: *A meme is a unit of cultural information that evolves by means of natural selection (i.e., selection, replication and variation), and which exhibits the replicator properties of longevity, fecundity and copying-fidelity.*

Having arrived at this definition, the next stage is to investigate what can constitute *a unit of cultural information* in music. Jan's definition of a meme in music as 'a discrete "packet" of musical information, demarcated from neighbouring material' (2007, p. 3) is used as a basis to argue that a unit of information can include a continuous monophonic group of notes. Miller's argument on how information is chunked into seven plus or minus two events in short-term memory (1956), and Jan's argument that a meme needs to consist of at least three notes (2007, pp. 60-61), is used to refine the group of monophonic notes to a contiguous series of three- to eleven-notes. These ideas on what constitutes a unit of cultural information are then merged with the definition of a meme to provide a working definition of a meme in music that can then be used to investigate memes using mass data analysis. A meme, for the purposes of this study, is

> *Any three to eleven monophonic consecutive notes, excluding symbolic ornamentation and secondary parameters, that evolves by means of natural selection (i.e., selection, replication and variation), and which exhibits the replicator properties of longevity, fecundity and copying-fidelity.*

The first chapter in Part II investigates how computers can be used to analyse music. A set of four factors for pattern finding using computers, based on Uitdenbogerd and Zobel's six factors (2004, p. 1054), is alighted upon: defining a pattern, defining similarity between patterns, sourcing a suitable encoding system, and sourcing a suitable music analysis program. The definition of a pattern from Part I covers the first factor. The second factor is more complex owing to the diverse nature of the disciplines surrounding the concept of similarity between patterns in music (ranging from motivic to structural, and from cognitive to analytical concerns). Additionally, investigating some of the solutions for solving the problems of determining similarity between patterns (i.e., neural networks, weighting systems, Levenshtein algorithms) is deemed to be beyond the scope of the present study. Therefore, a definition of similarity between patterns based on some of the musicological properties of patterns is implemented (the similarity algorithm).

For the third factor, a number of different encoding systems are investigated, ranging from audio- to score-based systems. The audio-based systems are discounted owing to the difficulties involved in analysing the audio files. Two score-based systems are alighted upon that hold the necessary information to conduct the present research: MusicXML and Kern. MusicXML is chosen over Kern due to its integration with programs such as *Sibelius* and *Photoscore*, as well as its compatibility with relational databases.

The final factor involved looking at music analysis programs such as Melodic Match and the Humdrum Toolkit. Both of these programs are deemed unsuitable: Melodic Match because of its handling of the definition of a pattern in music, and the Humdrum Toolkit due to the complex nature of how it deals with patterns compared to alternative technologies. An alternative approach of using a relational database is then investigated. It is found that the relational database will work with the present research because it is designed for mass data storage and analysis, has excellent indexing capabilities, has a standard querying language in the form of SQL, and is used as part of Knowledge Discovery in Databases (a concept developed to aid mass data analysis).

The next chapter in Part II (Chapter 4) then describes the stages of Knowledge Discovery in Databases and shows how the concept can work for music analysis. This involves showing how the music was selected, encoded into MusicXML, and transferred into a relational database model. Following on from this it is then shown how the data is stored, enhanced and transformed within the relational database in order to compare the frequency with which patterns appear across time, using a ranking system. The chapter concludes by explaining two different types of testing undertaken. Firstly, testing took place to check that the algorithms were working by using samples of data and manually checking the results. Secondly, a set of pseudo-random notes were generated and passed through the algorithms together with a set of pseudo-random notes seeded with three pre-determined patterns to check the methodology.

The final chapter of Part II looks at the initial results of the live data, explaining some of the consequences of the methodology, such as the predominance of same-note pattern types, and the resultant use of a de-duplication algorithm.

In Part III, the ranking positions of patterns are analysed for evidence of selection, replication and variation as stated in the working definition of a meme in music. The argument advanced is that if a pattern becomes more frequent across time, then it is being replicated and selected more often than a pattern that becomes less frequent across time. Because the ranking positions show the relative frequencies of patterns across time, they are used to show that certain patterns are becoming more, or less, frequent than other patterns, and are consequently showing evidence for selection and replication.

Variation is more difficult to find evidence for because it involves determining an ancestral link between patterns. A similarity algorithm is used to show that there are possible links between some of the patterns that first appear in the middle period used within the dataset, and patterns that have appeared in earlier periods. It is argued that because there are new patterns in the middle period that have possible antecedents in the previous periods, variation has taken place. Some of these new middle-period patterns are then shown to become more frequent over the later periods than other patterns. Therefore, there are new patterns in the middle period that have possible antecedents, which then show selection and replication over the remaining periods.

Evidence for the replicator properties of longevity, fecundity and copying-fidelity is not sought directly from the ranking positions and similarity algorithm. It is argued that it is the meme instance that exhibits these properties. Therefore it is not possible to produce direct evidence for these replicator properties from the dataset because it only holds information on one instance (a single printed edition of a score) of one manifestation (the score) of a meme. However, Dawkins argues that the replicator properties determine the abundance of some memes over others in the meme pool (1989, pp. 194-195). His argument is used to suggest it is possible to infer that the replicator properties exist because there is evidence from the ranking positions of patterns over time that some patterns are more prominent than others. When the figures are investigated further by examining the abundance of patterns across compositions as well as within the pool of patterns, the evidence for abundance becomes less convincing. Additionally, when using the similarity algorithm to investigate copying-fidelity, it is found that the data does not always behave as would be expected.

Part IV summarises the issues considered during research design, the chosen methodology and the investigations undertaken. It is important to be clear that the present study does not attempt to offer statistical modelling of memes and their properties. Statistical modelling is appropriate when there is a large population of objects of the same sort with a distinct categorisation, and the goal is to draw general conclusions about their properties by selecting a small sample for study. Statistical modelling ensures that the generalisations made from the sample are valid for the whole population. In this research, there is a large population of patterns *some* of which may exhibit distinctive and unusual properties that set them apart from the larger population; i.e., are possible memes. The goal is to find such individual patterns, and not to draw conclusions about all patterns within the population.

The present study offers an investigation of one approach to searching for memes in music, together with a number of suggestions for improvements and enhancements. The result of the investigation is that some evidence is found for potential memes in music based on the corpus of works analysed, the working definitions used, and assumptions made within this study.

Part I – Literature Review

1 Chapter 1: What is a Meme?

1.1 Introduction

As explained by Wilson (2009) towards the end of 2007, a person (or possibly persons) sat down and wrote a list of personal facts and posted them on the social networking site *Facebook*.[1] This idea was taken up by others and was passed around in the style of a chain letter within *Facebook*. The recipient was supposed to write down sixteen facts about themselves and pass these onto sixteen friends together with the original instructions. According to Wilson, this then evolved as the number of facts was changed in the instructions by the participants. Eventually, the number settled at twenty-five and then the idea took off rapidly before experiencing an equally rapid decline (Wilson, 2009). This demonstrates an idea that has unknown origins, that then mutates with a number of different variations with one variant becoming dominant, before the idea declines rapidly. Wilson (2009) suggests that this idea is a 'meme', but what does this actually mean?

Dawkins (1989) introduced the concept of the meme based on the idea that cultural and evolutionary concepts can be combined. However, since the introduction of the concept of the meme a number of commentators have produced a wide variety of definitions for the meme. This chapter begins with Dawkins' concept of the meme and how the idea has spawned a number of differing definitions. It argues that it is difficult to produce a definition of a meme that can accommodate a number of conflicting ideas on what constitutes culture, as well as a definition that can provide answers on how memes are stored and manifest themselves.

1.2 The Original Concept

In his book *The Selfish Gene* (1989), Dawkins looks into how and why genes replicate and, on this basis, introduces the idea of a cultural equivalent to the gene. Dawkins draws from his work on genes to formulate the idea that genetic evolution is not the only form of evolution. He argues that cultural evolution exists and draws an analogy between it and genetic evolution. Dawkins believes that there is a cultural equivalent of the gene with similar properties and actions. With this analogy, he is

[1] www.facebook.com

abandoning the idea that culture is purely for the benefit of genetic replication, arguing that culture has a separate evolutionary existence from genes (1989, p. 191). Dawkins uses the term *meme* for this cultural equivalent to the gene (1989, p. 192).

To illustrate the concept of the meme, Dawkins talks about memes being 'tunes, ideas, catch-phrases, clothes fashions, ways of making pots or of building arches' (1989, p. 192). These ideas spread themselves through societies by 'leaping from brain to brain via a process which, in a broad sense, can be called imitation' (Dawkins, 1989, p. 192). To illustrate this, Dawkins uses an example of scientists hearing about a hypothesis, deciding whether it is acceptable or not, then if it is, passing it on to colleagues and students, which results in the spreading of the idea (1989, p. 192).

This idea of a meme has, to a certain extent, caught on with a small number of academics. For example, commentators such as Blackmore (1999) and Dennett (1999) have been promoting the concept of the meme amongst both the academic community and the wider community as a whole. *But what exactly is a meme?*

1.3 Defining a Meme

When Dawkins introduced the term 'meme' in 1976, the concept and analogy had not been fully formulated (Dawkins himself expanded on the analogy later on (1982, pp. 109-112)). Unfortunately this has led to several different accounts of what a meme is. Wilkins claimed more than twenty years later that the term was still unclear in what it represents (Wilkins, 1998). Plotkin, two years on from Wilkins, pointed to the importance of getting definitions correct; contested definitions, he argued, can lead to researchers going in different directions, causing problems with the coherence of terms (2000, p. 73). The implication from both Wilkins and Plotkin is that clarity is needed on what constitutes a meme otherwise the use of the term will lack credibility. Aunger points to another problem with defining a meme by saying that as many writers on memetics are ignorant of advances in anthropological research, they will inevitably make mistakes that have already been made when dealing with the evolution of culture (2000, p. 224). He makes the claim that 'no one knows what a meme is' (2002, p. 21) and puts this down to there being no proof of the existence of memes.

Dawkins, in his later book (1982), gives a more coherent definition of a meme by saying that it is 'a unit of information residing in a brain. It has a definite structure, realized in whatever physical medium the brain uses for storing information' (1982, p. 109). Dawkins qualifies this by saying that a meme's 'phenotypic' effect is the physical manifestation of the meme, using clothes and music as an example (1982, p. 109). This defintion places the storage of memes in the brain but does not say how the memes are replicated or how they evolve. However, Dawkins does go on to explain that memes are copied by their phenotypic manifestations imprinting themselves into other brains, but he does not go into detail of how this is achieved (1982, p. 109).

Aunger gives a definition of a meme that follows Dawkins' idea that memes are stored in the brain by saying that memes are

> essentially the state of a node in a neuronal network capable of generating a copy of itself in either the same or a different neuronal network, without being destroyed in the process. (2002, pp. 325-326)

Here Aunger is clearly arguing that memes are stored in the brain, as suggested by Dawkins, by introducing 'neuronal networks' into the definition. He further suggests that replication takes place by inducing copies in other brains. There are two problems with this definition. Firstly Aunger, like Dawkins, does not make it clear how one meme induces a copy in another brain. Secondly, identifying memes with neuronal networks does not make the search for memes any easier because their effect on the brain can only be observed by using brain scans, which is extremely time-consuming and costly.

A similar approach to Aunger's definition of memes is taken by Blackmore, who describes memes as 'instructions for carrying out behaviour, stored in brains (or other objects) and passed on by imitation' (1999, p. 43) . This definition is less technical than Aunger's but the idea of memes being stored somehow in the brain is the same in both Aunger's and Blackmore's definitions. However, Blackmore also places ' other objects' in the definition as an alternative means of storing memes. With this definition, Blackmore is separating the meme from the storage mechanism and giving a means whereby memes can replicate. The important point for Blackmore is that memes are passed on by imitation and uses Dawkins and the *Oxford English Dictionary* to justify this assertion (1999, p. 43). Thus, for Blackmore, memes cannot exist without the ability to replicate through imitation.

Brodie describes memes as

> a unit of information in a mind whose existence influences events such that more copies of itself get created in other minds - this allows for yawns not to be memes but the opening of Beethoven's 5th to be one. (1996, p. 32)

Like Blackmore, Brodie is emphasising where memes are stored and how they replicate. Unlike Blackmore, who places the storage of memes in both the brain and other objects, Brodie places the storage of memes firmly in the brain, just like Aunger and Dawkins. Brodie also goes further than

Blackmore by beginning the definition with a view of memes as 'a unit of information' as opposed to just 'instructions'. The problem here is that Brodie does not define what constitutes the 'unit of information' or how such units are copied.

Borenstein's definition also suffers from the problem of what constitutes 'information'. For him,

> Memetics is an evolutionary theory of culture, one based on the competition among ideas and bits of information for the attention of our brains, which will perpetuate them by repeating them. (2004, p. 465)

This links in well with Brodie's definition but expands it by including the idea of competition, selection and replication, which is also missing from the definitions of Dawkins, Aunger and Blackmore. An additional problem with Borenstein's definition is that the storage mechanism for memes is only implied.

Wilkins defines the meme as 'the least unit of sociocultural information relative to a selection process that has favourable or unfavourable selection bias that exceeds its endogenous tendency to change' (1998). Like Borenstein, Wilkins is bringing in the idea of memes being subjected to evolutionary forces by talking about selection and change. Wilkins also goes further than both Brodie and Borenstein by describing the information as 'sociocultural'. Whilst addressing the idea that memes evolve as well as the problems of defining what type of information they consist of, Wilkins' definition avoids stipulating where memes are stored.

Gabora gives a very general definition of a meme in a title to a section in an article: 'Meme: the Unit of Information that Evolves through Culture' (1997). Unlike other commentators such as Aunger, Dawkins, Blackmore and Brodie, Gabora is not stating where memes are stored or that they are replicated. In addition, Gabora's definition does not define what constitutes a unit of information, how it evolves or what makes up culture. Gabora's definition has an advantage over some others in that it allows the search for memes to be carried out by observing their manifestations in a culture. Also, like the other definitions, Gabora's definition places culture as being the only evolutionary force acting on memes. There are other forces that can act on the evolution of memes such as biological and environmental forces. For example, the structure of the human larynx will have an impact on some musical memes; tunes that are more easily sung may become more dominant. Additionally, the weather can impact on the likelihood of open-air music making within a culture (music making in the rain is not that pleasant), with different memes being required for open-air than indoor music making in order to ensure the sound carries over a wider distance. Also, culture cannot exist without information, but information can exist without culture (for example, a gene is a form of information that exists independently of culture). Under Gabora's definition, non-cultural units of information could also evolve through culture. In order to clarify the above point, perhaps a better definition would be:

A meme is a unit of cultural information that evolves by means of natural selection.

However, this definition of a meme still has some problems. Like Borenstein's, it does not state explicitly where memes are stored. This can be an advantage as it does not restrict the search for memes to specific areas. It also does not specifically address how memes are replicated, however it could be argued that replication is a component of natural selection. The definition also assumes knowledge of how evolution operates and what constitutes culture.

1.4 The Vagaries of Culture

In *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, Dennett claims that the only aspect of life that makes man different from other animals is that man has evolved culture (1995, p. 338). White also makes this observation by arguing that 'Man is Unique: he is the only living species that has a culture' (1959, p. 3). Gabora also makes this claim by asking 'Why is Culture Unique to Humans? - A Speculative Answer' (1997). In this article Gabora suggests that only humans have evolved culture because only humans can cope with complex thoughts owing to the enhanced vocal and manual abilities (1997). However, defining culture is difficult because different disciplines approach the concept from diverse angles.

Back in 1949, T. S. Eliot wrote his *Notes Towards the Definition of Culture* (1949) in which he looks at the questions of whether religion is part of culture, and whether culture is determined by education. The title is interesting by itself in that it implies he is not giving a fully formed argument by using the term *notes*, that these *notes* do not relate to an exact definition of culture by using the qualifier *towards*, and that he thinks that there can be an exact definition by using the definite article. Unfortunately, nearly sixty years later, the *Dictionary of Sociology* states that 'there has not unfortunately been much precision in its use' (Abercrombie, et al., 2006, pp. 92-93).

Eliot identifies three conditions for culture (he accepts there are more but concentrates on these): firstly, there needs to be an 'organic structure' that will allow for culture to be passed between different generations; secondly, that it should be possible to monitor different cultures in different regions and determine their differences; and thirdly, that there should be a ‘balance of unity and diversity in religion’ (1949, p. 13). It should be noted that for the first two he is talking about culture, and for the third he is talking about religion (here he is firmly placing religion as a part of culture). Eliot recognises that there are different areas to culture such as art, scholarship and philosophy, and that someone excelling in one of these areas may not excel in the others (1949, p. 21). Additionally, Eliot argues that as society becomes more complex, different cultural levels emerge across societies and he asks whether these are socially hereditary or whether the cultural levels propagate themselves (1949, p. 23). Eliot also states 'that culture is the one thing that we cannot deliberately aim at. It is the product of a variety of more or less harmonious activities, each pursued for its own sake' (1949, pp. 17-18), which leads to his tentative definition of culture: 'Culture may even be described simply as that which makes life worth living' (1949, p. 26). This definition fits in with Dennett's assertion that culture is what makes man different from animals. However, it does not take into account that there are people with culture who think life is not worth living but do so due to the social and legal pressures against suicide. These pressures are themselves part of our culture.

Whilst Eliot was contemplating definitions of culture, anthropologists were working in this area as well. Aunger explains that the anthropological definition of culture has evolved over time, beginning with the idea that culture has to do with objects, social contracts and religious beliefs (2002, p. 30).

An example of this type of definition is provided by White, who sees culture as 'an extrasomatic, temporal continuum of things and events dependent on symboling' (1959, p. 3). This definition identifies the importance of culture as its ability to give additional meaning to an object or idea. White cites the use of holy water as an example, saying that it is only man that can distinguish this from ordinary water and that animals can not appreciate such subtleties (1959, p. 4).

Aunger argues that this type of definition (i.e., culture has to do with objects, social contracts and religious beliefs) is rather a 'catch-all' approach to culture which does not give a broad enough idea of what is *not* included in culture (2002, p. 30). He then explains that cultural definitions moved (from the 1960s onwards) onto the idea that culture is something formed in the brain, as 'ideas, beliefs, and values', but rejects this, saying that there has been no real consensus to this type of definition (2002, p. 30). Nevertheless, Harris relies upon this second type of definition in saying that 'culture refers to the learned, socially acquired traditions of thought and behavior found in human societies' (1997, p. 88). He quantifies this by saying that when anthropologists are talking about culture, they are looking at the lifestyles of whole societies, including their thoughts and patterns of behaviour, not just their artistic and literary achievements (1997, p. 88).

Finally, Aunger gives his own definition as 'a collection of ideas, beliefs, and values that can be abstracted from individuals and considered as a pool of information at the population level' (2002, p. 30). Yet this definition only goes some way to meeting Aunger's own concerns about early definitions of culture in that it is rather a 'catch-all' definition and does not say what is not included in culture.

Cavalli-Sforza and Feldman define culture as

> the total pattern of human behavior and its products embodied in thought, speech, action and artefacts, and dependent upon man's capacity for learning and transmitting knowledge to succeeding generations. (1981, p. 3)

Mead uses a similar definition to the second half of Cavalli-Sforza and Feldman's definition in her preface to Benedict's book *Patterns of Culture* (1934, repr. 1989), saying that it is 'the systematic body of learned behavior which is transmitted from parents to children' (Benedict, 1934, repr. 1989, p. xi). Although Mead doesn't define 'learned behavior', she does qualify the transmission, albeit rather narrowly, by limiting it to the parent/child line. These definitions go further than Aunger's in that they include the idea of *transmission* rather than just *abstraction*.

Aunger's and Cavalli-Sforza and Feldman's definitions extend Dennett's assertion that culture is what differentiates man from animals in that they include reference to *belief*; and Cavalli-Sforza and Feldman include reference to *artefacts*, the existence of neither cultural trait having been sufficiently proven in animals. However, Mead's definition does not help Dennett's assertion that culture differentiates man from animals in that it only refers to 'learned behaviour', and this is something that has been demonstrated in animals.

Sociologists approach the concept of culture differently from anthropologists. For example, the *Dictionary of Sociology* lists six approaches to a concept of culture (Abercrombie, et al., 2006, pp. 92-93):

1. culture contrasted with biology by comparing with non-biological manifestations of societies;
2. culture contrasted with nature by showing that humans are civilised and not barbarians;
3. culture contrasted with structural groupings through institutions and organisations within societies;
4. culture contrasted with materials using beliefs and ideas;
5. culture as a way of life with differences in the way societies behave; or
6. culture as a means of distinguishing high and low culture.

However, most writers on culture only offer a definition in one or two of these areas and do not address all six.

Durham approaches a definition of culture by doing a survey of the literature on what constitutes a unit of culture (1991, p. 188). The common theme he found was that a unit of culture is an 'idea' but he argues that this is too general and lacking in precision. Similar words such as 'belief' and 'thought' are also rejected on the same grounds. He comes up with five properties that culture (as opposed to a unit of culture) must have (Durham, 1991, pp. 3-7). These are:

1 - 'Conceptual Reality': culture should have shared values and beliefs,
2 - 'Social Transmission': culture should be learnt from others,
3 - 'Symbolic Encoding': culture should give meaning to symbols,
4 - 'Systematic Organisation': culture takes on its own structure, and
5 - 'Social History': culture should be traceable between generations.

Yet these properties by themselves do not give a clear indication as to what a unit of culture is, only how it manifests itself and how it behaves.

Plotkin argues that culture is the ability to learn about the ideas and learning of others (1994, p. 213), which matches some of the properties suggested by Durham. He regards culture as a product of human intelligence, defining it as the ability to cause changes to our behaviour (Plotkin, 2000, p. 70). However, humans are able to train other animals to modify their behaviours, such as the toilet training of cats and dogs. This modifying of animal behaviour goes against Dennett's assertion earlier in this section that culture is what makes man different from animals. However, Plotkin does admit that the boundary between human and animal culture is fuzzy and introduces the term 'protoculture' when talking about chimpanzees displaying behaviour similar to human culture (2000, pp. 72-73). Whilst Boyd and Richerson agree that some animals display some signs of culture, they argue that this is the result of the ease with which the animals can learn certain behaviours independently of the actual transmission of culture (2000, p. 151). According to Plotkin, there are three levels of intelligence. Firstly there is 'learning and memory', secondly there is 'reasoning and thought' and thirdly there is 'shared learning and knowledge' (1994, p. 154). He suggests that each level involves a smaller group of animals with the last level having the smallest group. This is presumably Plotkin's way of allowing for humans to be different from animals.

Ellen argues that anthropologists have struggled to find a definition of culture that allows it to be explained in terms of sufficiently independent units that correspond with the idea that the brain can digest, separate, and reform units into new meanings (2002, p. 6). This he argues is due to the tendency of anthropologists to talk about culture from an evolutionary and/or diffusionist stance. Anthropologists are not the only ones to think about culture in terms of evolution, as Aunger points out by saying that

> Social scientists have long remarked that the pool of beliefs and values held in common by members of social groups – their culture, in short – appears to evolve over time. (2002, p. 2)

Therefore, Ellen and Aunger respectively are saying that anthropologists and social scientists believe that culture can evolve.

According to Hull, anthropologists have been investigating cultural and conceptual change for a number of years and uses Richard Semon's book *Die Mneme als erhaltendes Prinzip im Wechsel des organischen Geschehens* written in 1904 as a case in point (Hull, 2000, p. 50). Bloch also gives a list of anthropologists (including Steward (1955), White (1959), and Levi-Strauss (Levi-Strauss, 1962)), who have all written about cultural change before Dawkins introduced the concept of the meme (Bloch, 2000, p. 191).

Boas (who is called the father of anthropology by Walker (2004, p. 159)) was working on the idea of cultural evolution during the first half of the twentieth century and his book *Race, Language and Culture* (1940) includes a number of essays on culture written during this period. Here, he not only looks at existing cultures, but at culture from a historical and developmental angle, with a chapter devoted to cultural evolution entitled 'Evolution of Diffusion' (1940, pp. 290-294) originally written in 1924. This chapter looks at folk-tales and family groups to show that culture can spread across different regions, and time-spans with subtle changes. Thus Boas is introducing the notion of cultural evolution well before Dawkins' inauguration of the meme. Similarly, Bloch points to Kroeber's work on culture (1952) being 'superorganic', arguing that this approach is similar to that of memetics. For Bloch, Kroeber is arguing that culture has reproductive mechanisms independent from biological reproductive systems (Bloch, 2000, p. 191).

Benedict talks about how one of the objects of anthropology is to understand the nature and differences of culture, together with how they manifest themselves, vary, and change across different peoples (1934, repr. 1989, pp. 1-2). Benedict argues that anthropologists ask if biological factors can impact on our cultural traits. She then argues that this is not always the case that biological factors impact our cultural traits by pointing out that North American Indian tribes are biologically similar but have different cultural behaviour (1934, repr. 1989, pp. 233-234). This goes against the idea that culture is just another function of genetic evolution. Additionally, Benedict makes a brief allusion to the importance of selection in culture, arguing that in order for there to be a coherent culture, a society cannot adopt all possible cultural variants and this must therefore lead to a selection process (1934, repr. 1989, p. 237).

White explores the idea that culture changes over time in *The Evolution of Culture: The Development of Civilization to the Fall of Rome* (1959). He states that 'man and culture originated simultaneously' (1959, p. 5) around one million years ago. This, White theorises, happened with the evolution of the ability to use symbolism in the brains of some anthropoids (1959, p. 6). White points out that culture is separate from biological ties by arguing that although babies are born with a brain that has the ability to use symbolism, it does not have a culture at birth. The culture is then acquired from the social environment around the baby (1959, p. 12). Here he is agreeing with Benedict's idea that biological factors do not influence cultural traits, and helps further explain why Benedict's Indian tribes have the same direct biological ancestors but possess different cultures. Additionally, White also points to the various strands of mathematical thinking that can be traced over time which creates an evolutionary picture of the subject as an example of cultural evolution (1959, pp. 30-31). White also argues that the first humans, compared with other animals, were not well suited to their environment. They had lost their covering of hair that gave warmth, could no longer escape into the tops of trees away from predators, were not particularly strong, and did not have offensive weaponry like stings and venom to attack other animals (1959, p. 77). Thus man's ability to survive must have partly come from culture, which White argues is 'a mechanism whose function is to make life secure and continuous for groups and individuals of the human species' (1959, p. 78). By saying this, White is implying that culture has developed from the human facility for survival but that culture is not ingrained biologically, because man is born without culture. However, this supports the argument that culture is there for the benefit of genetic evolution because White is saying that without culture, man could not have survived.

The examples above from Boas, Benedict and White show anthropologists talking about cultural evolution before Dawkins introduced the concept of the meme. Both Benedict and White argue against a direct causal link between cultural behaviour and biological evolution, although White argues that culture may have originated due to the inadequate physical attributes of humans. Additionally, Benedict admits that there are selection forces on culture and White shows that cultural lineages can be compiled. However, none of the definitions or writings on cultural evolution points to Ellen's (2002, p. 6) and Dawkins' (1989, p. 192) idea that it is the splitting up of culture into discrete units which can be manipulated by the brain that is the key to a definition of culture.

The fact that there are many alternative views on culture highlights that there is no agreement on a definition of culture. Looking at various writings on cultural evolution has also failed to provide any consensus as to what constitutes culture and cultural evolution. This does not help in expanding the definition of the meme as *a unit of cultural information that evolves by means of natural selection*.

1.5 Storage, Manifestation, and replication of Memes

It is difficult to imagine a situation where culture could exist without the presence of animals shaping and propagating it. So what role do animals, and more specifically man, play within culture, and how is culture represented, stored and passed between generations? Memeticists, such as Blackmore and Dawkins, agree that man plays a role in the transmission and storage mechanisms, but there are variations between commentator's views on how this works, and how all the processes fit together.

When originally expounding his meme concept, Dawkins identifies the brain as the place for storing memes. As part of this, he argues that this is also where selection takes place because there are limited resources in the brain so it will choose which memes to hold (1989, p. 197). This idea is picked up by Dennett, who also believes that the brain stores memes and that the limited supply and size of brains causes selection pressures (1999, p. 131). This also means, according to Dennett, that memes need to have the power to get themselves selected over other memes (1993, p. 206). In Dawkins' later book, *The Extended Phenotype: The Gene as the Unit of Selection* (1982), he goes further by arguing that some memes will only be successful in an individual's brain if they tie in with other similar memes already residing in that brain (a bit like confirmation bias). He illustrates this point be saying that memes relating to communism will do better when they are transferred to a brain dominated with a set of memes that already relate to communism (1982, p. 111). Conte makes a similar point by saying that memes are spread by social influence, which implies that successful memes need to conform to current social thinking in order to be successful. Additionally, Conte takes the idea that similar memes are attracted to one another, and adds a further dimension by arguing that some memes interfere with other memes. Conte uses the example of a car being flashed by another car coming towards it arguing that the driver of the car being flashed originally takes the flashing as a signal of greeting. The initial interpretation of a greeting signal is then changed to one of warning when the driver sees a speed control ahead (2000, pp. 109-110). Thus the initial greeting meme formed from an association with other greeting memes within the brain is transformed into a warning meme. Another point is made by both Durham (1991, p. 210) and Milicevic (1998, p. 27), who argue that the workings of the brain are not all conscious with Durham (1991, p. 210) also believing that this can lead to unconscious selection.

If all this selection is going on in the brain, does this mean that the human brain behaves differently from animal brains? Human brains are different from other species, according to Dennett, because they have 'habits and methods, mind-tools and information, drawn from millions of other brains which are not ancestral to our own brains' (1995, p. 381). But does this account for the relatively large size of the human brain?

Plotkin points out that brain size in humans is about seven times what it should be in relation to the body (1994, p. 54). According to Aunger, the brain has increased in size due to needing better organised and more efficient memory, together with the necessary backup mechanisms to counteract the fact that neurons die (2002, p. 213). He also makes the point that if you view the increase in terms of Darwinian evolution, bigger brains must have evolved because they were better suited to the environment (2002, p. 182). However, this does not explain what in the environment caused brains to become bigger.

Calvin also makes a similar point about humans having a number of unused neurons at birth which die out later on, but goes on to say that brains continue to develop into adulthood (1997). Milicevic also points to the human brain developing after birth, with it being initially only around twenty-five percent of its eventual weight (1998, p. 27). The fact that the brain is developing after birth also ties in with Dennett, who says that there are certain variations in the brain between people that get fixed for life after birth, for example learning to speak a language when young fixes certain patterns of neurons within the brain which are different for all individuals (1993, p. 183). Such issues regarding the development of the brain are further complicated by the concept that the brain is able to refine its own structures and functions, i.e., its plasticity (Kolb & Whisha, 1998).

Aunger goes onto argue that the change from the ape brain to the homo brain was due to a change in the structures of the brain and not in how it processes information. However, he contradicts his own argument to a certain extent by saying that this led to new cognitive abilities that allow planning, memory, conscious thought and communication skills, all of which can be said to involve the processing of information (2002, p. 231). In addition, he neglects to point out that these abilities are found in various animals (such as apes, as shown by Darwin (1871 repr. 1981)), who do not have this enlarged brain.

There is also the problem that human memory is rather a haphazard storage mechanism. Calvin points to the fact that there is a lot of guesswork taking place in memory. The brain will use previous memories to fill in gaps in knowledge, quite often incorrectly, and he cites the example of constantly overlooking a misprint in a book as an example of this behaviour (Calvin, 1996, p. 39). Plotkin also makes this point by saying that memory is constantly reorganising itself and, as a consequence, changing its perception of what is known (1994, p. 180). He also brings in the idea of guesswork but from a different angle by saying that the brain will often make guesses rather than invest in the effort of working out complex calculations (1994, p. 192). Both Lynch (1998, pp. 10-11) and Milicevic (1998, p. 28) pick up on the inaccuracies of memory by saying that the brain fuses together past experiences and recombines them to generate new ideas and altered versions of the original experience. Milicevic sums this up by saying 'our memories, as exact, recorded, fixed images of the past, are an illusion' (1998, p. 28).

An experiment conducted by Sperber illustrates this point. He took two different diagrams and asked a group of people to copy them one after the other, like a visual *Chinese Whispers* game. The first diagram did not contain recognisable shapes but by the end of the process it had mutated into recognised shapes with a clear path of how it had changed. The second diagram contained recognisable shapes which were more or less copied accurately (so that they were recognisable as a copy of the original) without any trace of a lineage between the original diagram and the end product. Here Sperber argues that the participants in the first diagram tried to create a mental image of the shapes that corresponded with shapes that they already knew and that this is what caused the mutation to conventional shapes (Sperber, 2000). This experiment shows that the brain is imperfect in creating copies of ideas which allows for the variation of ideas, something that Blackmore argues for by saying that variations are caused by imperfections in the memory as well as imperfections in communication (2000, p. 28).

Blackmore also takes Dawkins' idea of memes competing within the brain and introduces the idea that genes also have an impact on meme selection. Her reasoning is that memes are competing with each other which affects the genes (such as the size of the brain), whilst the genes are also trying to select the most beneficial memes for gene replication (for example, imitating the ‘best hunting technique’ memes will increase the chance of genetic replication) (Blackmore, 2000, p. 37). Blackmore introduces a new theory with the concept of 'memetic drive' (2000, pp. 36-37). Memetic drive is driven by imitation, i.e., memes are copied from one person to another by imitating each other. Once imitation of memes had begun, people then needed to decide which memes to imitate. At the same time, those genes that allow for successful imitation helped to reinforce the imitation of memes. Therefore, memes that help the survival of genes (tool-making etc.) are allowing genes to help with the survival of memes through better imitation, i.e., a positive feedback effect. This ultimately ends up changing the brain, with Blackmore arguing that ‘in this way, memes force genes to create a brain that is capable of selecting from the currently successful memes’ (2000, p. 37).

Others disagree with Blackmore's ultimate assertion that imitation is the only mechanism for replication. Conte accepts that memes can spread through imitation but includes other types of transmission as well, such as monitoring of behaviours and social learning (2000, p. 98). Plotkin also agrees that imitation is used to replicate memes but adds that language and learning are also used for the replication of memes as they both help with the copying of information (2000, p. 76). Language for Plotkin is a means of gaining further knowledge which, he argues has a central role in culture (2000, p. 83). However, Blackmore does not deny the role of language, for she says that language is better at spreading memes than visual communication (2000, p. 33). However, Blackmore believes that language is a means for imitation and thus replication.

In order for replication to take place, there must be some form of transmission, whether by imitation or other means. Blute believes that there are two aspects to the transmission of memes. Firstly, there is virus-like transmission, which is where memes spread horizontally through society regardless of generational boundaries. Secondly, there is the gene-like transmission, which equates to vertical transmission where memes are passed between generations (Blute, 2005).

Lynch advocates virus-like horizontal transmission by arguing that 'actively contagious ideas are now called memes' (1998, p. 3). Brodie also advocates virus-like transmission of memes by saying 'your thoughts are not always your own original ideas. You catch thoughts - you get infected with them, both directly from other people and indirectly from viruses of the mind' (1996, p. 14). According to Brodie the mind is good at dealing with instructions and is also adept at learning, which makes it more susceptible to virus-like memes (1996, p. 63). Brodie gives three ways in which the mind can be infected with memes. The first is by 'conditioning', which happens through learning, repetition and association. The second is 'cognitive dissonance', where the mind seeks solutions for opposing information by tying itself to certain sets of similar memes. Finally, there are 'trojan horses', whereby a meme sneaks into the brain by attaching itself to other memes (Brodie, 1996, pp. 138-143). However, these three methods for memes to get into the mind can be explained in terms other than infection. 'Conditioning' can be explained as learning from the environment, be it by Blackmore's imitation (1999, p. 3) or through other forms of learning such as trial and error. Both 'cognitive dissonance' and 'trojan horses' can be explained by Dawkins' (1982, p. 111) and Conte's (2000, p. 109) idea of the memes being both attracted to, and more acceptable, within an associated grouping of memes (i.e., confirmation bias). Laland and Odling-Smee also point out that viruses depend on the susceptibility of potential hosts in order to propagate (2000, p. 134). If this is applied to memes, then this accords with Dawkins' and Conte's idea for if the host already agrees with the communist ideology, then the susceptibility to communist memes will be greater than the susceptibility to capitalist memes.

According to Cavalli-Sforza and Feldman, there are three forms of transmission: vertical (i.e., from parents to children in the same way in which genes are passed on), horizontal (i.e., peer to peer), and oblique (similar to vertical but not from direct relatives) (1981, p. 54). They divide these forms various further into modes (1981, pp. 55-60):

1. Vertical
   - Parent to offspring
   - Other family members in same generation as parents to child
   - More remote family generations to child
2. Horizontal
   - Siblings
   - Age peers
3. Oblique
   - Non-biological family to child
   - Members of social groups other than family
   - Teacher to pupil
   - Society notables
   - Media
   - Groups and societies

Within these modes, there are differences in what is transmitted. For example, it is argued that mothers will influence their children with classical music and horoscopes, that fathers will influence their children with sports and parties, and that the maternal/paternal mix will influence their children with religion and politics (Cavalli-Sforza & Feldman, 1981, p. 89). All of these influences are determined by Cavalli-Sforza and Feldman using statistical models, but they accept it is difficult to get accurate statistics owing to the mix of other influences involved, such as the child's peers and social groups (1981, p. 89). Ultimately, they use statistical modelling to show how similar the spread and evolution of culture is to epidemiology.

Dennett suggests that memes are 'parasitic' rather than viral (1993, p. 200). However, some of his arguments give the impression that memes are virus-like. He talks about memes slowly getting to a certain level of saturation before spreading rapidly to become generally accepted. This is demonstrated by using the example of publisher's best-sellers lists that often give a book far more exposure than books not in the best-sellers lists, and consequently dramatically increase the sales of best-sellers (Dennett, 1999, p. 133). Additionally, Dennett looks at the speed memes spread and says that they 'now spread around the world at the speed of light, and replicate at rates that make even fruit flies and yeast cells look glacial in comparison' (Dennett, 1993, p. 203). In contrast, Hull points to the fact that genes can also be passed on rapidly, using the example of viruses and bacteria, and he concludes that both meme and gene propagation can be fast moving (2000, p. 55).

Blackmore argues that if memes have only vertical transmission, then they are all transmitted alongside genes and that this means there would be little or no conflict between the two (1999, p. 132). From this, she deduces that there must be horizontal transmission (in which she includes Cavalli-Sforza and Feldman's oblique transmission) as well, and goes on to say that horizontal transmission works independently of gene transmission because there is no direct connection between the two (Blackmore, 1999, p. 133). Additionally, Blackmore points out that in order for memes to become totally independent of genes, horizontal transmission needs to become the major mode of transmission. She argues that this is now happening with memes transmitted by parents becoming less influential as more memes are being transmitted by other sources such as 'schools, radio, television, newspapers' etc. (1999, p. 134).

Aunger argues that memeticists should be looking at the idea that memes are replicators in order to understand them (2002, p. 3). He discusses how symbols can be used to transfer memes, but argues that such symbols are not complex enough to hold all the information that is required for replication to take place. From this Aunger decides that memes cannot just jump from brain to brain and that there must be some other mechanism (2002, pp. 236-237). In his words, 'memes can't be translated from brain stuff to signal stuff and back again' (2002, p. 236). To try and overcome this problem, he looks at the way prions replicate by using a protein to pass on information , and the way that computer viruses link themselves to other programs to replicate (2002, p. 103) (however, computer viruses can have more complexity than just linking themselves to existing programs and they do not all act in this way). From the discussion of symbols, he decides that something extra is required and calls this an 'instigator' (2002, pp. 240-241). Instigators are used by memes to aid their replication in a similar way that prions use proteins to pass on information and computer viruses use other programs to propagate. In other words, the memes stay in the brain and the instigators act as a means to transfer the information. The main problem with this reasoning is that he does not make clear how the idea of the instigator can enhance his idea that symbols are not sufficient to convey memes. In other words, this does not appear to progress the thinking on how memes are transferred.

In addition to his idea of an instigator, Aunger also sets out some conditions that are required for replication. These are:

> Causation: The source must be causally involved in the production of the copy – Similarity: The copy must be like its source in relevant respects – Information transfer: The process that generates the copy must obtain the information that makes the copy similar to its source from the same source, and – Duplication: During the process, one entity must give rise to two (or more). (Aunger, 2002, pp. 73-74)

Do these conditions fit in with Blackmore's imitation? For imitation to take place there needs to be at least two persons taking part; firstly the person undertaking the original action (the originator), and secondly, the person imitating (the imitator). If the imitator copies the originator then the originator is 'causally involved': without the originator, the imitator cannot imitate. There must be a point at which, if the imitator is digressing too far from the imitation of the originator, then the originator’s and imitator’s actions are different things and the process breaks down. This is similar to the problem of deciding at what stage a meme becomes a new meme rather than a variation of the original meme when copied (as will be discussed in Section 2.3 below). Before that point the imitator must be at least roughly copying the originator, thus meeting Aunger’s second condition. Likewise, the same argument can be used for the third condition, information transfer, in that if the imitator is copying the originator, information transfer is taking place. Once the imitator has copied the originator, then both the imitator and the originator have copies of the meme, which meets the fourth condition, duplication.

Sperber also stipulates some conditions for replication,

> 1 - B must be caused by A, 2 - B must be similar in relevant respects to A, and 3 - The process that generates B must obtain the information that makes B similar to A from A. (2000, p. 169)

These can be mapped onto Aunger's conditions, with the first matching Aunger's causation, the second matching similarity and the third matching information transfer. The only difference is that Aunger explicitly states that there should be duplication whereas this is only implied by Sperber. As such, the same arguments concerning whether Blackmore’s process of imitation satisfies Aunger's conditions apply to Sperber's.

So far, the discussion has concentrated on memes in terms of the brain, memory and imitation. However, the artefact also has a role to play in a theory of memes as either a possible container for memes or a possible transmitter of memes. Aunger argues that many artefacts exhibit evolution by the fact that they show signs of being adapted and becoming more complex (2002, p. 278). Plotkin claims that culture can be stored exosomatically in artefacts but goes on to say that this does not mean that culture did not exist before artefacts were around (1994, p. 214). If artefacts are a form of storage mechanism for memes, what does this mean for replication?

Boyd and Richerson say that artefacts hold information but that this information is minimal. An example is a clay pot which holds information on what it is and some of its properties, but does not show how the pot can be made (Boyd & Richerson, 2000, p. 147). This will of course create difficulties for replication in that a pot is not easy to reproduce without supplementary knowledge. Sterelny argues that it is not ideas that are copied, but artefacts and skills. He says that in order for artefacts to be successfully copied (in other words, replicated), the materials to make a copy need to be easily available, the artefacts need to allow reverse engineering to take place, and there needs to be a certain amount of leeway for errors (Sterelny, 2006, pp. 156-157). Whilst it is fairly easy to copy a book because there is a large amount of technology to help someone to do this, it would be considerably more difficult to make a copy of a mobile phone. Both of these artefacts have evolved and are in abundance, and could be said to be cultural products. However, there is the problem separating out the physical being of an artefact and the information that it holds. Books not only instantiate the properties of a book (for example a cover, a number of pages with writing, made out of paper, etc.), they also include information on how it was made (from paper and printed on), as well as holding the text which is information in itself. A book can also hold information on how to replicate a book, or some other cultural artefact, for example a recipe for a pudding. Aunger would argue that the meme in a book is the content of the text. He notes that when the story of Don Quixote is transferred to different media, the story is still the same so it is the 'information content that defines a replicator, not its material embodiment' (Aunger, 2002, p. 156), thus separating out the medium from the message. However, he neglects the point that the medium itself could be a meme (books have evolved over time from being hand-written, through the advent of printing, to the use of e-books) and as such the medium is also open to replication.

If the view is that memes are units of information, then because a book stores information, the book must be able to store memes. But if books store memes, how does replication take place? Books are mass-produced and therefore a large number of copies are made, but does this copying constitute replication according to Aunger's conditions? Aunger himself takes an example of an artefact and tests it against his replication conditions. For his test, he uses the Queen's head on a stamp. He argues that it meets the first criteria of causation because the Queen signed the decree for the stamps to be manufactured, the second of similarity because the picture is a good likeness, the third of information transfer because it is the Queen's head on the stamp, and finally of duplication because a number of copies of the stamps are produced (Aunger, 2002, pp. 75-76). As for books in terms of their text, the author causes the book to be written and then copied, the copy contains the information of the original, the copy has got its information from the original and finally, the book is a duplicate of the original. This means that a book's text can meet Aunger's conditions for replication, and as such it could be argued that books store memes. If this is accepted then where does this leave the idea that memes are stored in the brain and passed on by imitation, as suggested by commentators such as Dawkins and Blackmore?

There clearly seems to be some disagreement on issues such as the storage, transmission and replication of memes, and the role of artefacts and man in memes. This disagreement is reflected in the various different definitions of the meme, which again highlights the need for a definition that does not go into detail about the storage, transmission and replication processes of memes until agreement is reached in these areas.

1.6 Memes versus Genes

If the brain is important to memes, how does this affect genes? Aunger cites an interesting example of conjoined twins. They had virtually identical genotypes and obviously the environment they grew up in was the same. However, one of the twins was more dominant than the other. Aunger uses this example to claim that genes do not control everything in life (2002, p. 183). So if this is true, what effect does this have on the link between cultural and biological evolution?

Dawkins argues genes are totally independent of memes but that memes require genes for propagation (1982, p. 110). However, for Dawkins this does not mean that the success of a meme relies on the success of a gene. For this he points to the idea of the suicide martyr whose genes will not get a chance to propagate after death; but the martyrdom meme may infect others and cause them to do the same (Dawkins, 1982, pp. 110-111).

Wilkins uses Shakespeare's *Romeo and Juliet* as an example of the distinction between biological and cultural entities. Juliet is interested in the physical being of Romeo and not his identity (and consequently his culture). For Juliet, Romeo 'being a Montague, and being the person she loved, were two distinct states' (Wilkins, 1998). He goes on to say that although there are two dimensions, biological and cultural, culture will not necessarily make an individual more or less adapted to the environment but the two can influence each other (Wilkins, 1998). Durham takes the stance that although genetic and cultural evolution are distinct, they both have an impact on human behaviour

(1991, p. 159). He goes further by arguing that both genetic and cultural evolutions influence each other and uses the term 'coevolution' to represent this phenomenon. Both Wilkins and Durham are in effect saying that the influence runs both ways between genes and memes.

But where does this all leave the meme/gene analogy and can it move the concept of cultural evolution forward? Many see problems not only in the analogy itself (such as Dougherty (2001, p. 89) and Kuper (2000, p. 185)), but also in the fact that genes are studied from a physical science perspective whereas memes are not (such as Bloch (2000, p. 189) and Hull (2000, pp. 45-47)).

Dougherty is cautious of the analogy. He subscribes to the view that there is no physical form to a meme in contrast with the gene. From this, he argues that there is no scientific method which can be used to prove the existence of a meme, which he believes is a problem with the analogy (Dougherty, 2001, p. 89). Kuper goes further than Dougherty and calls the analogy 'fanciful and flawed' (2000, p. 187). Also Kuper argues that if memes are ideas and techniques, then it is difficult to separate them and treat them in isolation in the same way one can with genes. Additionally, he argues that the mechanism for transmission of ideas is very different from the transmission of genes (Kuper, 2000, p. 187). He uses these reasons to say that memes are 'shadowy' and only get shape because of the analogy. For these reasons he thinks that it is dangerous to use the analogy as it may cause differences and inconsistencies between the genes and memes to be ignored (Kuper, 2000, p. 185).

Bloch argues that in order for the meme/gene analogy to work, theories are needed that will work with both concepts. This is difficult because memeticists tend to come from a sociology background whereas geneticists come from a biology background (Bloch, 2000, p. 189). Others, such as Hull, have a problem with the cross-disciplinary nature of memetics. He makes the same point as Bloch but adds that each discipline seems to think that the other disciplines are simpler than they actually are (2000, p. 45). Kuper also criticises Dawkins for not citing any literature on anthropology and because of this, believes that Dawkins has formed a concept that does not add anything to the cultural evolution debate (2000, p. 179).

There is another problem with the analogy in that the definition of the gene is not secure. Both Wilkins (1998) and Hull (2000, p. 47) point to the fact that the definition of a gene has not been consistent, and that molecular biology is constantly changing its views on genetics. Wilkins also looks at the idea that genes by themselves are not always evolutionary by pointing out that a son is not a clone of his father as they can have many different physical characteristics (1998). However, this difference in physical characteristics could be regarded as just a part of the evolutionary process, for only when differences occur can there be a selection process to decide which differences are advantageous and beneficial to the gene-pool.

1.7 Summary

Cultural evolution exists in some form or another. Borenstein (2004) shows how the Soviet Union tried to control cultural evolution before the collapse of communism, and how western cultural ideas spread and gained influence in post-Soviet Russia, adapting and building on existing Russian culture. This example shows cultural evolution at work. In addition, anthropologists such as Boas, Benedict and White had been exploring the idea well before Dawkins and have given a number of examples of cultural evolution. So the real question here is can Dawkins' meme concept add anything to the debate on cultural evolution?

In order to answer this question, a clear definition of culture and what is meant by evolution is required. Darwin (Darwin, 1859 repr. 1985) produced a theory of evolution which is now generally accepted amongst the scientific community. However, it is difficult to arrive at an acceptable and comprehensive definition of culture. This makes forming a definition of a meme even more difficult. So far, no-one has been able to prove the existence of memes to the general acceptance of the academic community. Until the existence of memes is proved, then surely an accurate definition of a meme cannot be agreed upon. As such, a 'catch-all' definition will suffice to point intrepid meme-hunters in the right direction. Consequently, a variant of Gabora's definition was suggested: *A meme is a unit of cultural information that evolves by means of natural selection*.

2 Chapter 2: Defining a Meme in Music

2.1 Introduction

Chapter 1: What is a Meme?*,* Section 1.3 provided a definition of a meme that is rather broad in its reach. This definition makes searching for a meme difficult as it does not give any real ideas as to the possible physical or abstract properties of a meme. Another way to search for memes is to look for evidence of the properties that a meme should exhibit. If an aspect of culture exhibits these properties then it could be evidence for a meme. Since a meme is a unit of cultural information that evolves (see page 26), the first aspect to investigate is what processes underlie evolution.

This chapter begins by looking at what properties a meme should exhibit, arguing that a meme should have both the evolutionary processes of selection, replication and variation (based on Darwin's theories of evolution by natural selection), as well as the replicator properties of longevity, fecundity and copying fidelity (as stated by Dawkins). Following on from these processes and properties, the chapter then tries to define a meme in music. This is done firstly by looking at possible storage and replication mechanisms in music, followed by ideas on how compositions are musically connected. Finally, Jan's (2007) definition of a meme in music is used as a basis to argue that a single line of consecutive notes can be considered as a potential meme in music, and a working definition of a meme in music is created from this hypothesis.

2.2 Processes in Evolution

In 1831, after graduating from Christ's College, Cambridge, Charles Darwin set out on his epic journey aboard the survey ship HMS Beagle after being recommended for the journey by the botanist Professor J. S. Henslow. On his return in 1836, he began work on a theory of evolution by natural selection which resulted, more than twenty years later, in *On the Origin of Species by means of Natural Selection,* first published in 1859. He outlines a theory of animal and plant life evolving over many millennia from the same origins, arguing that species have adapted to their environment by evolving characteristics that allowed for an increased chance of survival, which in turn increased reproductive success rates and therefore allowed the passing on of those characteristics. He argues that there are a number of checks and balances that are at work to stop a single species from taking over an environment with food supply, predators and the environment all playing their part (Darwin, 1859 repr. 1985, pp. 119-120). Other factors such as physical borders can act as barriers (such as islands) stopping the spread of some species (Darwin, 1859 repr. 1985, p. 131). All this could be described as a battle and Darwin himself reinforces this view by talking about the 'war of nature' (1859 repr. 1985, p. 129), where 'the vigorous, the healthy, and the happy survive and multiply' (1859 repr. 1985, p. 129).

Patterson summarises natural selection with three deductive arguments. Firstly, species produce far too many offspring in relation to the number that are needed for replacement in order to keep the numbers at a constant level, therefore there must be a struggle for survival. Secondly, species exhibit hereditary variation. If these variations are useful for survival then species with these variations are more likely to survive and pass on these variations to the next generation. Thirdly, the environment is constantly changing, thus species whose hereditary variations help cope with the environment are more likely to survive (Patterson, 1987, p. 237).

Darwin's theory could therefore be viewed as a circle going between a species that adapts to its environment (variation), to some of these adaptations being favoured according to their fitness within the environment (selection), to those that survive and pass on their adaptations to their offspring (replication), back to those offspring adapting to their environment etc. (Figure 2.1). But does Figure 2.1 actually show the essential properties that are required for evolution to take place?

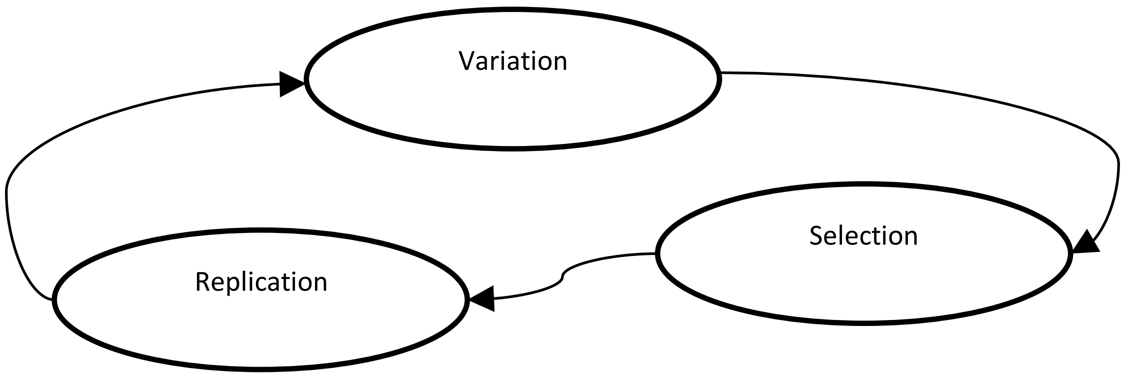

Figure 2.1: A circular view of the properties for evolution by natural selection to take place

A number of commentators have their own views on what is required for Darwinian evolution to take place. These vary in the number of processes involved, with some having just three and others having up to six, as well as in the terminology used. Aunger opts for three processes, consisting of heredity (the passing on of attributes), variation (the changing of attributes) and fitness (which attributes ensure the survival of the species) which can arguably be mapped into Figure 2.1 with heredity equating to replication, and fitness equating to selection (only those that fit best into the environment are selected) (2002, p. 25). Dennett also opts for the three-process model along similar lines as Aunger, saying that evolution occurs with

> 1 - variation: there is a continuing abundance of different elements, 2 - heredity or replication: the elements have the capacity to create copies of replicas of themselves, 3 - differential “fitness”: the number of copies of an element that are created in a given time varies, depending on interactions between the features of that element and features of the environment in which it persists. (Dennett, 1995, p. 343)

Again these processes tie in with Figure 2.1, however Dennett's order is different in that the variants are replicated before they are selected.

Gabora suggests four processes are required for evolution to take place. There must be some form of information that can go through the process to begin with, then there must be a way of generating variations of this information, then there must be a means and rationale for selecting some of these variations, and finally, there must be a way to replicate or transmit these variations (Gabora, 1997). The major difference between this model and the models of Aunger and Dennett is that Gabora starts with some form of information (which is taken for granted in Aunger and Dennett). The other processes map easily onto the three-process models of Aunger and Dennett. Durham's model (created before Gabora's) has an additional process to Gabora's model. This is the ability to be able to isolate one piece of information from another piece of information that can ultimately lead to variations, because a piece of information in two different environments without the ability to communicate with one another can evolve differently (Durham, 1991, p. 22). This last point fits in with Darwin's idea about species evolving separately on different islands, but this could be viewed as just another means of generating variations. Calvin uses a six-process model. Four of Calvin's processes are the same as Gabora's, but Calvin adds the idea that there need to be different success rates for patterns, and that some of these patterns should not survive (Calvin, 1996, pp. 104-105). These last two processes are really just an expansion of the selection element of Figure 2.1 and thus do not add anything new to the models.

If it is accepted that Calvin's final two processes (i.e., different success rates and some patterns not surviving) of evolution are an expansion of the selection component, and that Gabora's first process of starting with some form of information is a given, as Aunger and Dennett seem to do, then this results in the three-process model as shown in Figure 2.1 of selection, replication and variation.

2.3 Properties for Meme Replication

When originally introducing the concept, Dawkins came up with three different replicator properties that are required of a meme and that fit in with the gene analogy: longevity, fecundity and copying-fidelity (1989, p. 193). In, *The Selfish Gene* (1989), Dawkins examines how the replication of molecules took place when the process initially began to take place. He hypothesised that in those early days there were some molecules that were more stable (i.e., less prone to splitting), and other molecules that split quite easily. Those molecules that survived the longest without splitting, i.e., which achieved longevity, became more abundant. Dawkins also considered the frequency of replication amongst molecules, i.e., their fecundity. Those molecules that replicated more frequently would become more abundant than those that replicated less frequently. Finally, Dawkins considered the accuracy of the replications, i.e., copying-fidelity. If a molecule had an error in one in ten of its replications, then it will be less abundant than if it had an error in one in a hundred replications (Dawkins, 1989, p. 17).

Later on in *The Selfish Gene*, Dawkins applies the replication properties hypothesised for molecules to the meme. Firstly, he argues that individual copies of memes must have longevity. This means that they should be able to last for a sufficient length of time without becoming immediately extinct and thus not being able to replicate and become abundant within the meme pool. However, he does not specify any time period for this or make any comparisons to the longevity of genes. Secondly, memes must have fecundity. In other words, they must be easily replicated otherwise they will become extinct. Again, the idea behind fecundity is to ensure that a meme becomes abundant within the meme pool. Dawkins does not go into detail about what makes a meme successful but uses the example of scientists digesting new ideas before deciding whether to accept them or not. The ideas need to be accepted to become memes (even if they are not true). Finally, memes need copying-fidelity. That is, they need to be copied with few or no errors in the copying. This will help the meme become successful because the more accurate copies it makes, the more likely it is to survive. Like longevity and fecundity, copying-fidelity relates to how a meme becomes abundant in the meme pool (Dawkins, 1989, pp. 193-194). Brodie points to the problems of copying-fidelity by saying that if it is too high, then nothing will change and therefore there will be no evolution; and if it is too low, then there cannot be any replication (1996, p. 67). Dawkins admits himself that he is not sure if memes do exhibit sufficiently high copying-fidelity as it is easy to alter memes through communication (1989, pp. 194-195) (as in the game *Chinese Whispers*).

But how do the replicator properties of longevity, fecundity and copying-fidelity accord with the evolutionary processes of selection, replication and variation? Longevity will allow a meme the chance to be selected and replicated. The longer it survives, the more likely it is to be replicated and passed on, giving it more chance to be selected. Fecundity also enhances a meme's chances of survival in that if it is abundant then there is more chance of it being selected. Finally, copying-fidelity is more difficult to make a case for in this model. Does it matter if a meme changes drastically if it gives the meme greater longevity and fecundity? There has to be some variation of memes otherwise there will not be any evolution. However, if the changes are extensive and lead to greater longevity and fecundity, then does it really matter to the original meme, because the original meme has still created successful descendants?

2.4 Refining the Definition of a Meme

In Section 1.3 above, the meme was defined as *a unit of cultural information that evolves by means of natural selection*. If it is accepted that evolution consists of selection, replication and variation, then the definition can be expanded to *a unit of cultural information that evolves by means of natural selection, (i.e., selection, replication and variation)*. Likewise, if Dawkins' replicator properties are also accepted, then the definition can be expanded further to *a unit of cultural information that evolves by means of natural selection, (i.e., selection, replication and variation), and which exhibits the replicator properties of longevity, fecundity and copying-fidelity*.

This extended definition raises a number of questions when it is applied to music, with the major one being: *what is the unit of cultural information in music*? Does the unit of information consist of a broad aspect of music, such as the style of the piece, or does it exist at the individual note level, or perhaps somewhere in between? Is it an amalgamation of all the elements of music, such as pitch, rhythm, dynamics, texture, structure, etc.? Could there be many different types of unit of information across the broad spectrum of music that are, either independently or in tandem, subject to evolutionary forces? And do these units of musical information exhibit the properties of memes?

In other words, can music fit in with the definition of a meme?

2.5 What is a Meme in Music?

Having created a definition of a meme, this definition can then be tested against music. Before this can be achieved, there are certain areas that need investigating in order to understand what can be a meme in music.

This section begins by looking at how memes in music can be stored and replicated. The discussion focuses on how information regarding music can be stored in media such as the brain, and artefacts such as CDs, scores and books, arguing that these media exhibit the evolutionary processes of selection, replication and variation, as well as the replicator properties of longevity, fecundity and copying-fidelity. Following on from this is a discussion of the impact concepts such as borrowing and allusion have on memes in music, in order to determine if these concepts can help show that memes exist. Finally, various definitions of a pattern in music are explored and their consequences for the present research are considered, which ultimately provides a definition of a pattern in music that can be amalgamated into a working definition of a meme in music for the purposes of this study.

2.5.1 Storage, Manifestation, and Replication of Music

A question that is raised when looking at what constitutes a meme in music is how cultural information in music, and hence memes in music, is stored. Section 1.5 above looked at how the brain and artefacts are both possible storage forms for memes. With music there is the added question regarding artefacts of whether the meme is stored in an auditory (such as mp3 format) or a visual form (such as graphical notation). This question can be related to broader ontological problems such as what is a musical piece: is it the performance or the notation that defines the piece (a question beyond the scope of this study)?

Exploring music in relation to the mind is an extremely large and complex area that covers a wide variety of concepts and ideas. It encompasses diverse areas such as cognitive and biological disciplines. Although it is impossible to do any of these disciplines justice in such a short space, there are some elements that can have an impact on the understanding of the concept memes that must be addressed. These elements include how information is transferred to, and retrieved from, the brain.

To begin with, there is the concept of chunking. Most commentators agree that information is stored in Short Term Memory (STM) using a system of 'chunking'. Brower defines a chunk as 'a subjectively defined unit based on prior learning' (1993, p. 21). Radvansky expands on this by saying that a chunk is a small group of connected pieces of information where each group has an overall meaning which is dependent on prior learning (i.e., there is a constraint on how chunks are formed based on previous experience). These groupings are then able to extend the capacity of the STM because chunks can vary in size and complexity according to prior learning. Miller gives an example of chunking and how it is dependent on prior learning by talking about telegraph operators who originally heard Morse Code as individual letters, but then gradually started hearing them as whole words, and then phrases, i.e., as chunks (1956, p. 81). This is an example of where prior learning and knowledge can increase the size of chunking in STM, thus expanding its capacity. There is general agreement that the capacity of STM has a limit of between 5 and 9 units of information, a chunk (at whatever level) being one such unit.

Information is then transferred from STM to Long Term Memory (LTM) by repetition: Brower (1993, p. 23), Radvansky (2006, pp. 150-151), and Huron (2006, pp. 229-130) all point to the importance of repetition in the process of storing information in LTM, with Brower also remarking that the amount of repetition required will be different depending on the individual involved (Brower, 1993, pp. 23-24). Brower also points to the fact that once information has been passed to LTM, it is more or less permanently stored and that this storage has an almost unlimited capacity (1993, p. 22). However the idea that memory has an unlimited capacity can contradict the idea that information competes for storage space within the brain, as discussed in Section 1.5.

If repetition is important for the process of passing information from STM to LTM then repetition in music could impact on the effectiveness of this process. Huron develops this theme by arguing that composers' repetition of musical patterns 'acts like an involuntary form of conscious memorization' (2006, p. 229) and goes on to say that it is unsurprising that the musical motif is short and memorable. The shortness of the musical motif also fits in with the concept of chunking in that short sequences of notes are grouped together in STM.

Memory is not just about how information is stored, but equally importantly how that information is recalled. Baddeley et al. point out that cues are required to trigger the retrieval of information and that the memories are then pieced together according to these cues (2009, pp. 165-180). According to Snyder, there are three types of cuing. Firstly, there is the 'recollection' cue where there is a deliberate attempt to recall a piece of information. Secondly, there is the 'reminder' cue where the cueing of one memory automatically prompts the memory of another. Finally there is the 'recognition' cue where an external stimulus automatically triggers the memory (Snyder, 2000, p. 70). However, recall of information is not perfect and Snyder goes on to point out some of the defects of recall, including the problem of interference where cues that are similar trigger the retrieval of unintended information (2000, p. 70). Radvansky also points to another problem with recall in that '(e)ven the act of remembering alters memory because the experience of remembering gets stored' (2006, p. 141).

The concepts of chunking in STM, repetition in LTM and cueing in recall can all have an impact on memes in music. For example, the length and make-up of the chunks in STM could influence the possible length of a meme in music. It could be that the ideal length of a meme in music needs to fit in with the ideal length of a chunk in STM in order for the meme to gain an advantage over other memes that do not conform to the ideal length of a chunk. Additionally, a meme that requires fewer repetitions to pass between STM and LTM could have an advantage over a meme that requires a greater number of repetitions due to the limitations on capacity and processing within the brain. Finally, if the cue system for recalling information is intrinsically inaccurate, then this could lead to errors in the recall of memes in music, leading to the variation of some memes. Those memes that are fairly simple (i.e., those which have a small number of notes that are predictable in terms of the movement between notes) may have an advantage in that their recall could be more accurate more of the time than more complex memes (those which have a larger number of notes with complex pitch and rhythmic patterns).

Human memory is not the only possible store for memes in music as it is also possible that memes are stored in artefacts, as argued in Section 1.5 above. Other means of storing musical information are recordings (an auditory record of a performance), notation systems (limited instructions on how to perform a piece of music), and descriptions of music (such as books on the history of music, compositional guides, and guides to the theory of music).

Both audio- and notation-based storage means have the advantage of being able to record complex musical detail. The interpretation and background detail, however, are subject to a listener's or reader's previous knowledge and experience. For example, the interpretation of a trill in music notation depends on knowledge of the performance practices at the time of the writing of the composition. Audio and notation storage means also provide a more permanent record of music and are subject to less variation from inaccuracy in storage and recall than is the case with memory, although mistakes in the performance or the edition will be replicated every time the performance or edition is reproduced.

Books on music, such as on its history or on compositional techniques, provide a different perspective on music. Although information from books cannot provide the detail on individual compositions that can be produced from recordings and scores, they do still provide information on music. Historical books can provide insights into a composer's compositional methods, cultural and historical attitudes, influences, and linkages to other contemporary musicians. Books on compositional techniques can also provide information on music. A compositional treatise can become influential and, as such, its ideas can be spread amongst many composers. The Italian *Partimenti* system was used as training for composers in Naples during the eighteenth-century (Sanguinetti, 2007, pp. 53-54) and was still being used to train composers in the late nineteenth-century such as Puccini (Baragwanath, 2011). This system underlies much of the music from Italy and as such was influential both because it was used for training purposes, but also in the spread of the compositions that were formed with a background knowledge of the *Partimenti* system.

As recordings, notation, and books all hold information on music then, if memes exist in music, these storage means should also contain memes. It can be argued that these storage means are reproduced widely (e.g., many copies of the same CD recording or score) and hence so are the memes that are held within them. However, do these storage means allow for the evolutionary processes of selection, replication and variation to take place, and do they also exhibit the replicator properties of longevity, fecundity and copying-fidelity?

Selection takes place in all three storage means. Recordings are selected by the listener based on a number of factors such as the listener's own likes, recommendations from others, randomly hearing something the listener likes, etc. Performers will also select editions based on a number of factors such as the perceived quality of the editorial judgements made in relation to interpreting the composer's intentions, the ease of reading, the cost, etc. Books are also selected by readers again based on a number of factors, including the reader's likes, recommendations, cost, etc. Replication also takes place within all three storage means. Recordings are copied both within their original formats (such as CD) and across other formats (such as CD to mp3), and scores and books are printed, and sometimes reprinted, producing many copies. Finally, variation can also take place within all three storage means. With recordings, there are many different interpretations of the same composition which can be influenced by previous performances and recordings, with scores there are a variety of editions of the same composition, and books with authors arguing with each other over aspects of history, performance and composition.

Jan argues that notation can provide longevity for memes in music, saying that '[f]or the longevity of musical memes, the greatest advance was the development of notation' (2007, p. 33) This, Jan argues, is due to the greater stability of the meme instance within notation (i.e., the printed form is less subject to alteration) than within the mind. Therefore the meme can survive in a particular form for longer in notation than in the mind (Jan, 2007, p. 33). The same can be true for both recordings and books regarding music because both are less susceptible to variation compared with versions stored in memory. According to Jan, fecundity is also increased by recordings and notation as both can be copied frequently without errors, thus enabling a growth in the circulation of the memes (2007, p. 33). When tackling copying-fidelity in the same section of the book, Jan does not refer to any of the storage means (2007, p. 34). However, recordings, notation and books regarding music will all allow high copying-fidelity regardless of the length of the meme due to the semi-permanent nature of the storage mechanism involved, and their mechanisms for ensuring accuracy.

The mind, recordings, scores, and books on music have all been shown to be possible receptacles for memes in music. However, none of these possible interpretations of how the storage means can allow for the evolutionary processes of selection, replication and variation, and the replicator properties of longevity, fecundity and copying-fidelity takes into account the importance of the composer in forming the original compositions.

2.5.2 Musical Connections and the Meme

Western European composers are (on the whole) the creators of pieces of music that are then either written down in notation or performed, or a mixture of both. There is a vast array of compositions in music, which raises the issue of whether there are any connections between these compositions in terms of helping to show possible evidence for memes in music. In other words, are composers through their compositions the means by which memes are selected, replicated and varied?

One of the ways by which composers can pass memes between each other is by borrowing music from other composers' compositions. Borrowing, as defined by Burkholder, is 'taking something from an existing piece of music and using it in a new piece' (1994, p. 863). The use of the word *something* is interesting, in that Burkholder is not specifically stating what is being passed between compositions. However, he does qualify the definition further by saying that the borrowing can include many aspects of a piece from its melody to its structure (Burkholder, 1994, p. 863). This clouds the search for borrowed material because once one moves away from specific areas such as melody, it becomes more difficult to provide conclusive proof of a connection. This difficulty in proving connections will also arise when searching for memes in music. Firstly, because there is no clear definition as to what constitutes a meme in music, it is difficult to determine where and what the connections between compositions should be; and, secondly, because the more the composer moves away from just quoting the melody (i.e., introduces variation), the more difficult it is to justify a link without other evidence such as historical evidence. This problem of providing a link between works is noted by both Gimbel (1989, p. 233) and Tarnawska-Kaczorowska (1998, pp. 74-76), who both use it as a requirement for musical quotation to have taken place.

Metzer believes that borrowing is not just about the melody, but also about the cultural identity of the piece being quoted (2003, p. 2). He goes on to say that there is a bond between both pieces involved in the quotation that provides a link between the different micro-cultures in which the pieces were written (Metzer, 2003, pp. 3-10). However, Metzer argues that if the listener does not recognise the material, then the cultural meaning and the link with the past is missing (Metzer, 2003, pp. 6-7). He goes on to say that 'if the original is not recognized at all, is there a quotation?' (Metzer, 2003, p. 12). However, if the function of the meme is to replicate and dominate the meme pool, then this problem of the listener not recognising the quotation in music does not apply as long as the meme is being replicated. It therefore may not matter to a meme whether a quotation has been made without the knowledge of the listener (or indeed, the composer).

The difficulties in proving a link between compositions using the concepts of quotation and borrowing when an exact match is not identifiable are also present when composers allude to other works and styles. According to Knapp, Brahms used a lot of allusions to the past (for example, in his Symphony no. 1, Brahms alludes to the key of Beethoven's Symphony no. 5) to create a connection to the past in his own style (Knapp, 1998, pp. 9-11). Korsyn also points to Brahms having a preoccupation with the past which is shown through his use of earlier forms and genres (1991, p. 53). With allusions it is more difficult to find specific connections between pieces than with quotations because allusions tend to cover a greater number of mixed elements , such as texture, chord progressions, rhythmic devices, etc. Having to trace all these different elements in music makes it difficult to find a hereditary line of selection, replication and variation that would help show the existence of a meme in music.

Another concept that reflects connections between compositions is the phenomenon of intertextuality. According to Klein, intertextuality concerns the connections that exist between all texts in existence. These connections are formed because when reading or composing a text, both the reader and writer bring previous knowledge of other texts into their understanding of the work. Ultimately Klein argues that all writing is full of references to other writings, be it from direct quotations to allusions, which can be done both consciously and subconsciously (Klein, 2005). Klein sums up this idea by using a quotation from Bloom (Klein, 2005, p. 1): 'The meaning of a poem can only be another poem' (Bloom, 1973, p. 94). However, as Korsyn points out, gaining the meaning of a text poses different problems from gaining meaning from a piece of music (1991, p. 43).

Korsyn (amongst others, such as Butler (2003) , Klein (2005), Koutsobina (2008) and Plasketes (2005)) tries to map the concept of intertextuality onto music. He argues that there are a great many similarities between some pieces (such as between Brahms’ Scherzo op. 4, bb. 329-332 and Chopin’s Scherzo op. 31, bb. 65-68 and Waltz op. 64 no. 2, bb. 1-4), that cannot be explained through quotations and allusions alone owing to the level and detail of information about the connection required, and that a model is needed to deal with such connections (Korsyn, 1991, pp. 4-5). That model, according to Korsyn, can be intertextuality. Korsyn then goes on to rephrase Bloom’s quotation above to read ‘The meaning of a composition can only be another composition’ (1991, p. 14). By using Bloom's *The Anxiety of Influence*, Korsyn expands the concept of intertextuality in music further by saying that in the same way that poets are struggling with their place in history (i.e., their relationship to poets past, present, and future), so are composers (1991, pp. 6-15).

An argument to support the concept of intertextuality in music could be derived from the concept of memes. Memes could account for the connections between works, such as the similarities between Lutoslawski's and Chopin's first Etudes, and Bach's Prelude in C Major from *Das Wohltemperierte Klavier,* identified by Klein as intertextual links. Klein argues that the similarities between these pieces include the position of the piece in the sets, the key, and the keyboard figurations (2005, pp. 4-10).

These similarities show a possible link between the three works which were composed at different times with different cultural and musical practices. If these similarities between the compositions are taken as memes, then there is a case that selection, replication and variation has taken place, especially as Klein points to the composers knowing the earlier works involved. However, the major problem with intertextuality is that it implies multiple links across a vast canon of musical works. Therefore in order to search for memes, the multiple links hypothesised by intertextuality would need to be identified and investigated to search for evidence for the evolutionary processes, as well as the replicator properties.

2.5.3 Defining a pattern in music

The concepts of quotation, allusion and intertextuality all show that there are possible links between compositions. However, none of these concepts helps in determining exactly what a meme is in music because the challenges involved in identifying the connections between the compositions are extremely complex. Therefore a less complex approach is required.

At the start of his book, Jan suggests a definition of a meme in music as

> a discrete "packet" of musical information, demarcated from neighbouring material by various kinds of articulation and consisting of a relatively small number of uni- or multi-parametric elements – a collection of pitches in a distinctive rhythmic garb. (Jan, 2007, p. 3)

Here Jan is arguing that memes are patterns in music, however patterns are not necessarily memes, for they will not all exhibit the evolutionary processes of selection, replication and variation and/or the replicator properties of longevity, fecundity and copying-fidelity. However, the definition is stating that the 'packet' needs to have boundaries that allow it to be separated from the surrounding material. If a pattern can be a meme in music, then the question of what exactly constitutes a pattern in music needs addressing.

At present, there is no consensus as to what constitutes a pattern in music. Even a simple melody by itself can create difficulties when determining a definition for a pattern in music. Rhythmic elements, phrasing, metre, inessential passing notes, etc. can all exist in simple melodies, and need to be taken into account when formulating a definition for a pattern in music. Even after formulating a definition applicable to a specific melody, the definition may not be appropriate for other melodies.

Other issues also need to be addressed when formulating a definition, including structural patterns (such as form), polyphonic patterns, harmonic patterns, rhythmic patterns, patterns from dynamics, overlapping patterns, nested patterns, etc. There are also more subjective areas to investigate, such as how patterns are perceived by a listener, and how to determine which patterns are significant to the listener or to the composer. This section will investigate some of these issues in relation to memes in music and will provide a working definition that can be used within the boundaries of the present research.

When looking for patterns, it is necessary to define the boundaries of the pattern. However, the boundaries of a pattern in music can be difficult to determine owing to a number of different factors, for example the length of the pattern. Meredith et al. highlight the problems of the length of patterns arguing that a pattern in music can range from a small motif to a large section (Meredith, et al., 2002, p. 322). Another difficulty arises as to whether ornamentation should be treated as part of the pattern. Additionally, there is also the problem that patterns can exist within larger patterns (for example, thematic patterns exist within structural patterns), a point raised by Lartillot (2005, pp. 379-380). Patterns existing within patterns (nested patterns) means that some patterns will begin before other patterns have finished, therefore it is difficult to determine which notes should be included in the different patterns.

An obvious clue to look at when deciding what notes to include in a pattern would be phrase markings, but patterns are not necessarily aligned with phrase markings, and a phrase may consist of a number of different patterns. Other secondary parameters, such as metre and dynamics etc., can prove useful in distinguishing patterns, but formulating a definition that encompasses all these different variables is a complex task. There is also the added difficulty that some of these secondary parameters can create patterns separate from pitches and rhythms; e.g., a crescendo mark immediately followed by a piano mark is often used and, as such, could be considered a pattern in itself.

Another element in determining the properties of a pattern is addressing the issue of what makes a pattern important. Most compositions will contain a basic pattern of the first three notes of a major scale in sequence, either as a continuous group of notes or spread across a larger set of notes. However, just because this pattern exits in these works does not necessarily mean that this sequence is an important pattern within them (i.e., the issue of prevalence). Conklin and Bergeron argue that due to the large number of patterns that can be found in a work or corpus, the patterns found need to be ranked in order of their importance (2008, p. 60). Lartillot makes a similar point by saying that relatively uninteresting patterns are often found but that the selection criteria used for finding patterns needs to be robust enough to pick up and filter all the interesting patterns (2004, p. 55). However, as Meredith et al. point out, deciding what constitutes a significant, important or interesting pattern is difficult because the attributes of patterns and their context can vary according to both patterns placement within a work and the work itself (2002, p. 322).

There are some methods for overcoming the problems of defining a pattern. The most restrictive method is to define a predetermined pattern, i.e., one that specifies all the pitches and durations of the pattern. For example, the definition could be a pattern such as: a quaver G followed by another quaver G followed by a third quaver G before a bar line, followed by a minim E flat, i.e., the opening motif of Beethoven's Symphony no. 5. The major drawback with using predetermined patterns is that all patterns that do not match the predetermined patterns will be ignored. The advantage of focusing on a predetermined pattern is that the problems of patterns within patterns, ornamentation difficulties, and problems with the significance of the pattern for the work can be reduced because the predetermined patterns can include or exclude such elements within the music.

Clifford et al. opt for this restrictive approach by using a set of predetermined patterns to test four different pattern-matching algorithms using lead sheet music. The use of both the lead sheet music (which consists of a melody line with chord symbols) and predetermined patterns allowed Clifford, et al. to simplify the process of effectively testing the different pattern matching algorithms because they were only interested in the melody (2005, pp. 311-317). However, in their conclusions, Clifford et al. point to the need for fast algorithms when searching large datasets, even when using predetermined patterns (2005, p. 317). Cope also opts for the restrictive approach by predetermining thirty phrases for pattern matching in order to help garner evidence for allusion within music. This, Cope argues, is on the basis that using a larger number of phrases than thirty would have impacted heavily on the time taken for the computer to process the information (2003, p. 19). As Clifford et al. and Cope demonstrate, the use of a restrictive definition allows for speedier processing of music information. However, the use of predetermined patterns is inappropriate for the present research project for two main reasons. Firstly, it is not known what constitutes a meme and therefore generating a predetermined pattern that fits a definition of a meme is not possible. Secondly, predetermined patterns will lead to other possible memes being ignored.

It is also possible to define a pattern as any two or more notes without any other restricting criteria (i.e., unrestrictive), such as distance apart both tonally and rhythmically, starting points, secondary parameters, etc. For example, in a piece with ten notes, one pattern from the piece can consist of the first and eighth notes. The main advantage to this method is that all patterns will be included. However, this method has a major drawback in that it will be very demanding on the computer’s processing power due to the large number of possible patterns that can be found. For example, a theme with just ten notes in it will have 1013 possible patterns under the unrestrictive criteria.[2] Additionally, although this method can address the problem of nested and overlapping patterns, it does not address the issues of whether to include ornamentation, or how to determine which patterns are important.

Having tried the restrictive approach first, Cope then opts for the more inclusive semi-restrictive approach by using the program Muse. Muse searches using every possible combination of consecutive notes to provide statistical analysis of the patterns within compositions (2009, p. 297). Haus, et al. also opt for the more inclusive definition of a pattern by starting with any two consecutive notes and then adding a single adjacent note each time until no more matches are found (2004, p. 1048). In other words, the definition of a pattern used in both these cases is: any two or more consecutive notes. However, this definition is not completely unrestricted, in that it only refers to consecutive notes.

The idea of using an unrestricted definition of a pattern in music is appealing when relating patterns to memes. There are a couple of important reasons for this. Firstly, using an unrestricted definition means that all patterns can be identified for investigation as potential memes. If all possible patterns are found, then statistical analysis can be used to determine whether any show evidence for the evolutionary processes of selection, replication and variation, and the replicator properties of longevity, fecundity and copying-fidelity. Secondly, using an unrestricted definition removes the problems of specifying nested and overlapping patterns, because all patterns will be found regardless of their context. However, the disadvantage is that it will take a large amount of processing time and produce a large number of results. Therefore a pattern definition needs to be found that will produce manageable and meaningful results.

Another factor to consider is that Dawkins' argues that the shorter the genetic unit, the less likelihood of the unit being broken up and, as a consequence, the more likely it is to survive in its original form (1989, p. 29). However, Dennett argues that a three-note melody is too short. He cites the example of the three-nucleotide amino acid arginine that propagates profusely but does not have enough of an 'individual' effect to be considered as a gene. Using this example, he then argues that three notes cannot have enough impact to make a melody (Dennett, 1993, p. 344).

The basis for restricting the number of notes within a pattern could be linked with how the brain handles auditory information, as alluded to by Jan (2007, p. 202). Miller argues that the brain processes information in groups of 'seven, plus or minus two events' (1956, p. 81). Therefore, the definition of a

---

[2] The figure of 1013 is worked out by using 2 to the power 10, minus 1, minus 10 (i.e., $2^{10} - 1 - 10$). This is because each note can either be included or not included, but then the pattern without any notes does not warrant inclusion, nor do those patterns with just one note.

pattern could be: five to nine consecutive notes. However, Jan (unlike Dennett) argues that a meme can consist of just three notes (2007, pp. 60-61). If the lower limit is extended to incorporate Jan's minimum length, then the upper limit should also be extended by the same amount to maintain symmetry around Miller's central point of seven. The definition would then change to three to eleven consecutive notes. This means that seven notes remains the middle of the range, and leads to the interesting question as to whether there are more patterns clustered around seven notes than around the extremes, possibly confirming Miller's hypothesis.

There are still problems with this definition. For example, it does not address the problems associated with ornamentation or secondary parameters. Neither does it give any indication of structural patterns (such as form) or more abstract patterns (such as style). However, omitting these elements at this stage offers a more manageable initial research programme. Therefore the working definition of a pattern in this dissertation will be:

Any three to eleven consecutive monophonic notes, excluding symbolic ornamentation[3] and secondary parameters.

2.5.4 Definition of a Meme in Music

Section 2.4 above gave a working definition of a meme as *a unit of cultural information that evolves by means of natural selection (i.e., selection, replication and variation) and which exhibits the replicator properties of longevity, fecundity and copying-fidelity.* Underlying this definition is the assumption that the meme/gene analogy is legitimate, i.e., that the evolutionary processes for genes are the same as those for memes, and that the gene replicator properties are the same as for the meme. In order to map this definition onto music, two further assumptions are required.

The first assumption is that music is part of culture. Although this might seem obvious, a consensus on what defines culture is lacking, as shown in Section 1.4 above. Without an accurate definition of culture, it cannot be shown that music accords with the definition. However, any definition of culture that does not include music is surely deficient.

The second assumption is that a pattern in music can be considered as a 'unit of cultural information'. As shown in the previous section, Jan considers a meme to be 'a collection of pitches in a distinctive rhythmic garb' (2007, p. 3). This definition encompasses the working definition of a pattern in music as

[3] E.g., excluding trills, mordents, acciaccaturas, etc., but including appoggiaturas.

any three to eleven monophonic consecutive notes, excluding symbolic ornamentation and secondary parameters. Although the working definition of a pattern will not include all the patterns that Jan is referring to, it does represent a tractable definition of a pattern. As such, this working definition of a pattern seems reasonable.

If all these assumptions are accepted, then it is possible to combine the definition of a meme with the definition of a pattern to provide a working definition for a meme in music for the present research. Therefore this study will consider a meme in music as:

Any three to eleven monophonic consecutive notes, excluding symbolic ornamentation and secondary parameters, that evolves by means of natural selection (i.e., selection, replication and variation), and which exhibits the replicator properties of longevity, fecundity and copying-fidelity.

2.6 Summary

The definition of a meme alighted upon in Chapter 1 was expanded to include the evolutionary processes and replicator properties that a meme needs to exhibit, changing the definition to: *A unit of cultural information that evolves by means of natural selection, (i.e., selection, replication and variation), and which exhibits the replicator properties of longevity, fecundity and copying-fidelity*.

Unfortunately this definition is still deficient when considering memes in relation to music. If music is a part of culture, then the problem of what is a 'unit of cultural information' in music needs addressing. Jan's definition of a meme in music was used as a basis to argue that a pattern of notes can be a 'unit of information', even if patterns of notes do not encompass all possible 'units of information' within music. From this standpoint, a working definition of a meme for the present research was developed: *Any three to eleven monophonic consecutive notes, excluding symbolic ornamentation and secondary parameters, that evolves by means of natural selection (i.e., selection, replication and variation), and which exhibits the replicator properties of longevity, fecundity and copying-fidelity.*

But does the phenomenon of *Facebook's* 'twenty-five things' introduced at the start of Chapter 1 qualify as a meme? The idea shows selection in that eventually twenty-five became the dominant number of facts. The idea was also replicated by being passed between different friends, and was varied in that the number of facts was altered when replicated. Thus it fits into the three-process model for evolution of selection, replication and variation. The idea also displays elements of fecundity, in that it did spread around the community, and copying fidelity, in that the majority of the instructions remained unchanged. Unfortunately, longevity is a bit of a problem in that once the idea had peaked, it rapidly declined in use. However, the idea was around long enough for it to spread and, as such, it can be argued that it did exhibit enough longevity to make it a meme.

Therefore, if *Facebook* has been able to produce a potential meme, can music?

Part II – Methodology

3 Chapter 3: Music Analysis using Computers

3.1 Introduction

According to Cope, the history of using computers to aid music analysis and composition commenced in the 1950s when work began on trying to encode music in a format that computers could process (2009, pp. 28-36). This led to the development of the MUSIC I-V system for encoding scores into a digital format. Thereafter, there were a number of contributions to the field. These included the Digital Alternative of Music Scores developed by Bauer-Mengelberg, Gould and Logemann's ALMA, Kassler's Intermediary Musical Language, and Robison's Musical Information Retrieval (Cope, 2009, p. 26). More recently, developments have included MIDI, Kern and MusicXML.

Cope also points to a number of different programs and techniques for manipulating the encoded music data that have been developed (2009, pp. 28-36). Youngblood used computers to analyse music investigating Markov Chains in Schubert, Schumann and Mendelssohn in the late 1950s. Gross used computers to count intervals in small individual pieces, and Laske looked at syntactical structures in the 1970s. Alphonce used computers to determine pitch-class sets and Blobach in the 1980s used statistical analysis on Bach chorales to determine whether the harmony was more dominant than the counterpoint. Ebcioglu created the CHORAL system in 1992 which also analysed Bach chorales by using a form of predicate calculus to try and determine any underlying fundamental rules in their composition. More recently, Huron developed the Humdrum Toolkit as a multi-purpose music analysis system, Pople created an add-on for Microsoft *Excel* to undertake harmonic analysis, Temperley and Sleator developed the Melisma Music Analyzer to derive properties from music such as key and metre, Mazzola and Zahorka developed Rubato to analyse music using a predetermined weighting system with a scientific basis, and Taube developed the Music Theory Workbench to undertake harmonic analysis of Bach chorales.

A number of these systems were created to analyse specific subsets of repertoire such as Bach chorales (Blobach, Ebcioglu and Taube). Others focused on particular properties of music such as harmonic analysis (Pople, Taube) or pitch-class sets (Alphonce). Only Huron's Humdrum Toolkit, from the systems mentioned above, was developed without a specific repertoire or musical property as the sole motivation of the system. However, other systems also exist that were not developed for specific repertoires, such as VExPAT (Santana, et al., 1998) and Melodic Match (Wheatland, n.d.). But are any of these systems of any use for the present research?

This chapter argues that four factors, based on Uitdenbogerd and Zobel's ideas, need addressing in order to use computers to search for patterns in music (2004, p. 1054): a definition of a pattern, a definition of similarity between patterns, a suitable electronic encoding system, and a suitable computer program. Because the first factor has already been addressed in Chapter 2 with a working definition of a meme in music, only the remaining three require further investigation. Of the three, defining similarity between patterns is the most problematic owing to the complex nature of determining similarity. Some of the issues surrounding the similarity of patterns are discussed before alighting on a definition based upon up to seven properties of a pattern: the number of notes, the overall shape of the pitches and/or durations, the pitch and/or duration centre, and the distance between the highest and lowest pitch and/or longest and shortest duration. Finally, the remaining two factors are discussed, arguing that MusicXML is an appropriate encoding system due to both the availability of encoded scores in MusicXML and its integration with programs such as *Sibelius*, and a Relational Database such as Microsoft's *SQLServer* would be appropriate due to its data-handling and manipulation capabilities.

3.2 Requirements for Computer-Aided Music Analysis

Uitdenbogerd and Zobel identify six factors that need addressing when using computers to query patterns in audio music files (2004, p. 1054). Firstly, music is often polyphonic which makes separating the different strands in audio files difficult. Secondly, music is often in different tonalities and the tonality of the music has to be taken into account when querying the file. Thirdly, there can be non-melodic notes that sometimes need to be ignored when matching patterns. Fourthly, a way to determine which line contains the melodic material needs to be found. Fifthly, if using a sung voice to create the searched for pattern, then the accuracy of the singing needs to be taken into account. Finally, audio files are not a series of notes (they are a set of audio frequencies) and consequently a way needs to be found to convert the file to a series of comparable pitches and durations.

Although Uitdenbogerd and Zobel state that the first factor is a consequence of using audio files owing to the difficulties in separating out the frequencies in to different instrumental lines (i.e., assigning the note frequencies to the appropriate instruments), there is a similar problem when using scores. It is possible to have more than one voice to a stave in music notation, therefore a means of separating out the different voices is still required. The next three factors that Uitdenbogerd and Zobel highlight are also relevant to using electronically encoded music notation. Music uses a number of different keys and if patterns are to be found between pieces with different keys, or within pieces with a number of different key centres, then a means of addressing how to deal with different keys is required. How to handle non-melodic notes also needs to be addressed, including deciding how important they are and when to ignore them if necessary. Uitdenbogerd and Zobel highlight the importance of extracting the melodic material in audio files and this has a direct comparison with electronically encoded music notation, in that there needs to be a method for determining which notes are melodic. The final two factors that Uitdenbogerd and Zobel highlight are specific to audio files and to their own methodology.

Consequently, Uitdenbogerd and Zobel's first four factors highlight some important aspects of using computers to analyse music whether using audio- or score-based systems. Music can be viewed as a series of interconnected patterns and finding and comparing these patterns is at present fundamental to using computers to analyse music. However, Uitdenbogerd and Zobels' first four factors point to difficulties with finding patterns, comparing patterns, and prioritising patterns. These four factors are significant for dealing with two problems. The first is how to define what constitutes a pattern in music, which relates to all four factors. The definition of a pattern in music will help to determine the form that a pattern will take, where extraneous notes fit in, and which patterns are significant. The second problem is how to define when two patterns are similar, which relates to the second, third, and fourth factors. An accurate definition is important to ensure that only patterns that should be considered the same or variants of each other are recognised as such.

There are two other factors that need addressing when using computers to analyse music. Firstly, which encoding system should be used? Although Uitdenbogerd and Zobel used an audio-based encoding system that fitted in with their research methodology (2004), there are also score-based systems such as Kern, as well as some encoding systems that can be used to produce both scores and audio, such as MusicXML. Secondly, which technologies and techniques should be used to undertake the analysis? This second factor encompasses a large number of computer programs, operating systems and databases that can all have an impact on the performance and complexity of the algorithms used to analyse the music.

Therefore, in order for computers to aid in the process of analysing patterns in music, four factors need to be considered. These are: defining the properties of patterns, defining similarity between patterns, sourcing or creating a suitable electronic encoding system for music, and sourcing or creating suitable computer programs.

Section 2.5.3 above gives a working definition of a pattern in music for the present research that can be used to address the first of the four factors regarding defining a pattern. Of the remaining three factors, the most problematic is how to create a suitable algorithm that will be able to create a link between possible variant patterns in music.

3.3 Defining Variant Patterns in Music

If a meme is a pattern in music, and if memes exhibit the evolutionary processes of selection, replication and variation, then evidence for the latter needs to be found. Selection and replication can be investigated by finding exact matches of patterns across time (see Section 6.2.1 below). However, variation cannot be investigated using this method. Variation between patterns can only be determined by looking at whether there is any form of relationship (i.e., similarity of pitches and rhythm) between patterns across different time periods. Therefore, before variation can be found, a judgement on what constitutes a relationship between patterns needs to be determined.

Lartillot's comment that 'a pattern discovery algorithm based on exact identity of pattern instances will not be able to recognize transformed patterns' (2004, p. 54) is relevant in this regard. He is pointing out that although certain patterns can be matched exactly there are others that can only be related, through altered notes or through other operations such as inversion. Additionally, both Meredith, et al. (2002, p. 322) and Lartillot (2005, p. 383) add to the problems of matching patterns by pointing to the difficulties of dealing with ornamentation and subsidiary notes, which can obscure the overall structure of the pattern; and Polansky (1996, pp. 291-292) points to secondary parameters, such as dynamics and articulation, having an impact on what constitutes similarity (even if patterns have exactly the same notes and durations). Therefore there are a number of different problems in trying to define the relationship between patterns in music. Additionally, the problem of defining the relationship between patterns is not just a matter of dealing with differences in the notes and secondary parameters. As Polansky points out, there is also the problem of the listener's perception, with different listeners having alternative ideas on what constitutes a relationship; Polansky goes onto to say that 'the perception that melody A is closer to melody B than it is to melody C implicitly assumes that not only is similarity defined, but also degree of similarity' (1996, p. 294). Another factor involved in determining similarity between patterns is that different measures of similarity for different genres and styles will produce varying results. Rizo, et al. suggest that 'no similarity measure performs the best for all tasks, genres, or music formats', and go on to state that 'to develop an effective musical similarity measure is anything but straightforward. There is no similarity measure that works well in all musical applications' (2011, p. 313).

A common solution when using computers to determine whether patterns are related in music is to implement weighting systems. This is done by giving values to various differences between patterns. A simple example of this is if there are two patterns, A and B, where one note differs by a semitone between the two patterns, then the difference between the patterns is given a weighting for pattern similarity of 1. If there were a second note that is three semitones different between pattern A and pattern B, then this second difference is given a weighting for pattern similarity of 3, making the weighting for the differences between the two patterns equal to 4. Cope, in his Sorcerer program, introduces a weighting system as a way to eliminate pattern matches that do not correspond metrically (i.e., making sure that strong/weak matches are eliminated) (2003, p. 24). Another example of a weighting system is found in the work of Rolland, who introduces an algorithm called FlExPat (Flexible Extraction of Patterns) that uses a weighting system to determine the relationship of patterns based on differences in length where a certain number of notes are the same (1999, pp. 334-350). In general, the weighting system approach relies on specifying a suitable metric that takes into account all possible significant variations in patterns as well as being able to recognise that the context in which the change appears can affect the significance of the variation. Additionally, there is still a decision to be made as to what value of the measure used should two patterns no longer be considered related.

However, there are problems in using a weighting system for the present research. As already stated, Rizo, et al. commented that is difficult to find a measure that sits comfortably across different styles of music. Additionally, the concept behind the weighting system is to provide a measure of similarity between patterns rather than an absolute categorisation of whether two patterns are similar. Determining a cut-off point to provide an absolute categorisation for each of the different musicological periods used in this research would require additional work beyond the scope of this research.

Another way to investigate how patterns are related is to employ the concept of neural networks. There are proponents of using neural networks in music analysis, such as Rabuñal and Dorado. They believe that neural networks can go further than just supplying data by arguing that neural networks can remove human input within some aspects of music analysis, such as style recognition (2008, p. 243). An example where neural networks have been used in music research is Sotiropoulos et al. who used neural networks to link similar audio files together to provide a 'recommendation agent' that highlights suitable alternative listening material (2008). However, neural networks, according to Cope, require complex algorithms and a great deal of training data (2004, pp. 12-13).

There are also a number of specific problems with using neural networks for the present research. Firstly, the majority of work on using neural networks has been on categorising elements, an area 'well suited' for neural networks (Haykin, 1999, p. 791). For example, a neural network should be able to take a single melody and determine the melody's tonality (for example, major or minor). However, the present research requires the comparison of two patterns to determine whether there are any similar elements between the patterns, rather than a simple categorisation. Secondly, in order for the neural network to be effective, a predetermined definition of similarity would need creating in order to generate a viable set of training data. The end result would then be subject to the accuracy of the definition and training data. Thirdly, different neural networks may have to be employed to match across different musicological periods. For example, the definition of similarity between two patterns is likely to be different between early and late period classical composers than late romantic and early twentieth century composers due to the differing nature of the periods and the stylistic constraints of the time. Fourthly, as stated by Bose and Liang, 'neural network algorithms and models are not inherently suited for implementation on general-purpose computers' (1996, p. 445). In particular, Bose and Liang note that neural networks are particularly suited for implementing in parallel programming environments, an advance topic in computer science. Finally, there is a wide range of neural network techniques and algorithms (e.g., Radial-Basis Functions, Support Vector Machines, Committee Machines, Principal Components Analysis, Self-Organising Maps, Information-Theoretic Models, Stochastic Machines, Neurodynamic Programming, Temporal Processing, etc.), all of which would need to be researched for their suitability in determining similarity between patterns in music. Therefore, an investigation of neural networks is beyond the scope of the present research.

An alternative solution would be to use the Earth-Mover's Distance metric (EMD) (a variant of the Levenshtein distance measure). According to Xiao, et al., EMD is 'a distance measure calculated based on the solution of the transportation problem, and the minimum cost necessary to transform one distribution to the other' (2011, p. 1846). The idea is to find the optimal solution required to move objects between two locations: in pattern matching terms, what are the minimum number of changes required to convert one pattern into another pattern? The appeal behind this approach is obvious, in that it provides a quantifiable link between patterns. However, in this context it is ultimately another form of weighting system, in that the distance between two patterns is determined by giving a numerical equivalent to the relationship based on the differences between the patterns. For the simple weighting example above, exactly the same calculation can be applied for the EMD value. An example of an EMD application in music is provided by Jan, who has used the EMD approach to match different patterns in terms of their memetic evolution (Jan, 2013). Therefore, the same problems highlighted in using a weighting system for this research also apply to the EMD approach.

Additionally, other academics have used a variety of different similarity measures on music. Rizo, et al. attempted to compare a number of different approaches to similarity measures (geometric similarity, bar-specific, quantized point-pattern similarity, tree similarity, n-gram similarity, graph similarity) using examples of polyphonic music (2011, pp. 314-316). Their conclusion was that the different measures should be combined, but that further work on a larger corpus of material is required (2011, p. 322).

The diverse nature of the disciplines surrounding the concept of similarity between patterns in music (from motivic to structural, and from cognitive to analytical concerns) means that this thesis cannot do full justice to the subject. Nor is it feasible to utilise one of the complex similarity algorithms, based variously on weighting systems, neural networks, metric distance, etc., to accurately encapsulate all possible variants, because the present research is not looking for specific pattern matches within a particular style or context, but for a more generalised system that will match across different styles and contexts. Additionally, researching and analysing the various different similarity measures is complex and time-consuming, as shown by Rizo, et al. (2011) and therefore beyond the scope of the present research. Consequently, the present research will use a musicological approach by implementing a definition for similarity between patterns based on some of the fundamental properties of musical patterns.

This study has provided a working definition of a pattern in music (see Section 2.5.3) that excludes the majority of ornamentation and all of the secondary parameters, leaving just the pitches and their duration. Therefore defining similar patterns will be done using these remaining elements, e.g., pitch and duration. This choice leads to a similar pattern algorithm which calculates a value for the number of notes in the pattern, the overall shape of the pitches and/or durations, the distance between the highest and lowest pitches, the ratio between the longest and shortest durations, the average pitch of the set of pitch classes within the pattern determined by the pitch boundaries, and the average length of the notes within the pattern (see Section 4.6.7 below). Only patterns that have an exact match on *all* of the calculated values stated are considered as similar. Therefore, in terms of the current research, similarity between patterns is defined as:

patterns that have an exact match on their calculated property values.

A number of other calculations were used in early trials of the algorithm, such as the distance of semitones between the first and last note of the pattern. However, these additional properties made the system too complex to produce meaningful data in a reasonable time frame. Although the final algorithm produces matches that might not be considered variants, and ignores some potential variants, the algorithm does produce a good number of matches where the two patterns are clearly similar (see Section 6.3.2). Thus, it is an effective first draft of an algorithm for identifying possible memes in music that can be expanded on with additional research.

3.4 Music Encoding Systems

The next factor that needs addressing is which encoding format should be used. In order for computers to analyse music it needs to be in a format that the computer can process. This has led to a number of different encoding formats for music being developed. Cope lists a number of formats which include MUSIC I-V developed in the 1950s and 1960s, MUSTRAN developed by Jerome Wenker in 1962, Digital Alternative of Music Scores (DARMS) developed by Stefan Bauer-Mengelberg in 1963, ALMA developed in the 1960's by Murray Gould and George Logemann, and the Intermediary Music Langauge (IML) developed by Michael Kassler, again in the 1960s (2009, p. 26). There are also a number of data-interchange formats for sharing music data between computer programs and some electronic musical instruments, such as MIDI and MusicXML.

Cope's list includes encoding solutions which have been developed for specific purposes and as such are difficult to use in other contexts. Selfridge-Field argues that 'every system makes sacrifices somewhere to optimize clarity of its preferred features' (1997, p. 5). Cope notes that the problem of individualised encoding systems has resulted in there being no common protocol for the encoding of music (2009, p. 26). Melucci and Orio also argue that the main encoding systems used today, such as Kern and MIDI, contain different pieces of musical information and none is a complete record of the music (2004, p. 1059). However, these encoding systems are developed for specific areas of music and computing, and consequently they reflect different aspects of music. Rolland recognises this point, citing that it is essential to choose an appropriate encoding system, otherwise the results will be impaired (1999, p. 336). When selecting an appropriate encoding system for the present research, there were two main issues to resolve. Firstly, does the encoding format hold the information required for pattern matching in an accessible way; and secondly, are there enough music data files using that encoding format to conduct the experiments with?

Selfridge-Field argues that, in terms of information encoded, there are three groups of encoding formats: sound-based, notation-based and analysis-based (1997, p. 28). However, there are problems with this categorisation. It does not include audio-based formats, for the sound-based encoding formats do not store actual sounds, only a coded equivalent of the sounds; and analysis-based is for analysing scores (i.e., is notation-based) *and* recordings (i.e., is audio-based) which therefore can be grouped with the sound- and/or notation-based encoding formats depending on the type of analysis. Therefore, this section will split encoding formats into two main groups: an audio-based approach designed to encode music performance, and a notation-based approach designed to replicate the workings of music notation.

Most audio-based encoding formats are designed purely for the reproduction of audio signals. For example, Pulse-Code Modulation (PCM) is used to store audio recordings that can be read by compact disc players. Generally the audio signal is converted by sampling the analogue signal using an analogue-to-digital converter (ADC) that is then managed by a codec (a program for compressing and decompressing data). These codecs can employ a variety of techniques to compress the file size, such as dynamic range compression, with some codecs being lossy (i.e., during the compression stage, some information is removed that cannot be retrieved or recalculated). When a given codec has been employed in creating the file, the same codec needs to be used for reading the file. This, therefore, has an impact on the type of software needed to read the file (see Wong (Wong, 2009, pp. 109-111; 180-181) and Burg (2009, pp. 263-265)).

The major disadvantage with using audio-based encoding formats for the analysis of music is that the information is sampled and quantised from the analogue signal, effectively storing a sinusoidal wave based on the frequency and amplitude of the changes in voltage over time (Burg, 2009, pp. 5-20). Sampling means that only information at certain points in time is used; for example, a sampling rate of 10Hz means that a sample is taken ten times a second. Quantisation is where the amplitude is rounded to a figure that can be held in storage according to the restrictions of the sample bit size. Holding the information in this way creates two main problems for analysing patterns in music. Firstly, in order to look for note patterns a system for separating out the different frequencies and instruments is required because at any one moment in the file the stored frequencies may represent a number of different notes and instruments. If the frequencies are not separated out into their separate notes and instruments, pattern matching can only take place at the global frequency level, meaning that the matching will not distinguish individual notes and instruments. Burg argues that 'it is nearly impossible to separate the frequencies of multiple instruments played simultaneously' (2009, p. 233), whilst Temperley argues that ‘deriving pitch information from actual sound ... is highly complex’ (2001, p. 11). Nevertheless, there are commercial programs available that can convert audio files into a score-based format, such as Neuratron’s *Audioscore Ultimate*. However, Figure 2.1 shows the start of Beethoven’s String Quartet op. 127, 1st movement, which clearly demonstrates the problem of matching notes to instruments by showing notes existing for the piano, cello, guitar, bass and voice. The second problem is that the audio signal does not give accurate details of some secondary parameters such as metre, textures, dynamics etc., nor is it able to distinguish analytical nuances such as enharmonically equivalent notes.

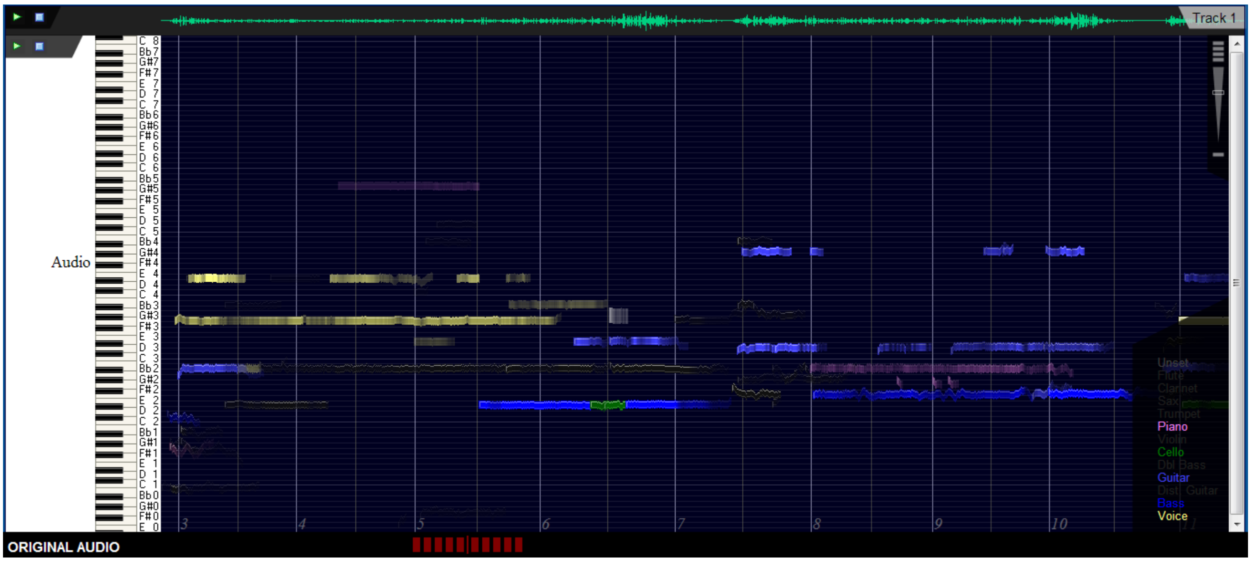

Figure 3.1: *Audioscore*'s analysis of Beethoven's string quartet op. 127 1st movement bb. 1-3

Nevertheless, the audio-based approach has a distinct advantage over the score-based approach in that it includes details of a performance that a score cannot include, such as the exact amount of rubato used or the precise grading between different dynamics. However, it must be remembered that some codecs use compression algorithms that remove certain pieces of information, such as inaudible frequencies, in order to reduce file sizes. For example, MP3 removes the top and bottom frequencies that the human ear cannot distinguish. Additionally, due to sampling, not all the audio signal is digitized, so care has to be taken to ensure that the same sampling rate is used across all files for comparison, otherwise there is the risk both of losing valuable information and comparing files that contain different levels of information.

Audio files are often used for research into music information retrieval. For example, Kostek has examined how to retrieve melodic phrases from audio files using techniques such as neural networks in order to create classification systems (2004, p. 1108). However, little research has been done on matching melodic patterns within audio-encodings, which could be due to the difficulty of trying to accurately separate out the different frequencies before melodic pattern matching can take place.

The second main type of encoding systems is those based on notation. Selfridge-Field describes a number of different notation-based encoding formats, developed over the years, designed to reproduce the equivalent of typeset music scores. These include MuTeX and its derivatives, SCORE, LIME Tilia, Nightingale Notelist, etc. (1997, pp. 31-32). Because of the emphasis on the graphical (i.e., where the notes are placed on the page) rather than notational aspects (i.e., what the note represents) of music scores, these encoding formats are not suitable for pattern matching analysis at the note level.

Therefore, this section will concentrate on three 'semantic' notation-based encoding formats: Kern, SMDL and MusicXML.

To begin, Kern is an encoding format developed by the Center for Computer Assisted Research in the Humanities as a format for the electronic distribution of music scores and is extensively used by the Humdrum Toolkit for music analysis (Huron, 2002). An overview of the Kern format is given by Huron, who describes Kern as a basic representation of western musical notation that concentrates on pitch, duration and instrumentation. Huron then expands on this description by listing six areas of notation that Kern encodes: pitch, duration, articulation and ornamentation, timbre, 'other' (i.e., performance indicators, etc.), and editorial markings (1997, p. 377). All these elements are given ASCII (American Standard Code for Information Interchange) text equivalents, which make them readable by humans as well as computers.

The advantage of using Kern for pattern matching in the present study is that it holds all the details required, i.e., pitch and duration. Additionally, there are a large number of Kern scores available online (Center for Computer Assisted Research in the Humanities, n.d.), removing some of the need to convert and import printed scores into a suitable encoding format. Using scores that have been created and uploaded by others means relying on the originators to ensure the quality and accuracy of the encoding. However, by using large numbers of scores, the impact of errors is reduced.

Another encoding system is the Standard Music Description Language (SMDL), which is the music-specific version of the Standard Generalized Markup Language (SGML) (Selfridge-Field, 1997, p. 36). SGML is, according to Prigmore, 'used to define different document models for different purposes' (2008, p. 15). Ultimately, SGML is designed to hold the actual data (i.e., the notes and durations, etc.), together with the structure of the data (i.e., how the different notes relate to each other), without any regard for any visual representation. Sloan lists four different domains that are used by SMDL to organise the musical information. Firstly, there is the *cantus,* which holds information that is relevant to both the score and the performance, i.e., the composer's name and the composition title. Secondly, the *visual* domain holds information about notation, i.e., which pitches follow one another. Thirdly, the *gestural* domain contains information about the actual sounds, i.e., the frequency of a pitch, and finally the *analytic* domain holds extra information and commentaries regarding any analysis of the piece, e.g., the details of the structure of the piece (Sloan, 1997, pp. 470-471).

Like Kern, SMDL is not held in binary but in text format. Additionally, SMDL pitch can be specified as either a note name or a specific frequency (Sloan, 1997, p. 472). However, SMDL does not offer any advantages over Kern in terms of pattern matching in the present research because they both hold all the information required, i.e., pitches and durations. The main difference is that there are a greater number of Kern scores easily available. There is also the problem of the complexity of SMDL, with Stewart saying that 'its approach is so abstract, unwieldy, and even obscure ... that a commercial implementation has yet to emerge after 15 years of availability' (2003, p. 60).

Finally, Extensible Markup Language (XML) is a simplified version of SGML and, according to Stewart, is ideal for representing music because it allows for flexibility in the way the document is structured (2003, p. 60). In the same way that SMDL is a version of SGML, MusicXML is a version of XML. Music XML was developed by Recordare, who based it on the MuseData and Kern formats, extending these to go beyond just western classical music. Unlike SMDL, a number of music applications have been developed that integrate with MusicXML, including the two main graphical notation software applications, *Sibelius* and *Finale* (Recordare, 2008).

Like Kern, MusicXML is not held in binary format, but uses ASCII text, making the files easier for humans to read (Recordare, 2008). However, the document structures can become extremely involved and complex owing to the large number of possible elements and attributes that can belong to an individual note, negating some of the benefit of being readable by humans. Again, MusicXML holds no advantage over either Kern or SMDL in terms of pattern matching in the present research because all three hold details of pitch and rhythm. Because MusicXML and Kern are related, it is relatively easy to convert between the two formats, which means that MusicXML and Kern are comparable regarding the number of scores available. However, MusicXML can hold more information about music than Kern, which could potentially slow down the performance of a computer because the MusicXML documents are correspondingly larger in size.

When referring to music encoding formats, reference also needs to be made to the widespread use of the Musical Instrument Digital Interface (MIDI) format. MIDI was developed as a common interface to allow communication between digital instruments and computers. Unlike audio files which store information about changes in voltage from analogue signals, MIDI stores information about the actual notes. This information includes details of the pitch, note length, instrument, loudness, etc. (Wong, 2009, p. 117). However, MIDI does not hold all the performance information that can be present in audio formats, leaving out information on performance nuances such as tone colour, etc.

There are disadvantages to using MIDI for matching patterns in the present research because of the way it encodes music. A given pitch has the same code regardless of any enharmonic equivalents, and duration is determined by *note on* and *note off* events. The former means that the subtlety of key, and moreover, the tonal nuances and implications, is missing from the files, and the latter must be translated into rhythms before analysing durational similarity.

As noted above, this study requires both an encoding system that holds all the necessary information, as well as a large sample of files that can be used to conduct the experiments. In order to use audio formats, work would need to take place to separate out the information required. Although there are programs to facilitate this (e.g., *Sibelius'* Audioscore Ultimate), these programs are not completely accurate and require a certain amount of user input to tidy the data. Notational formats such as Kern and MusicXML hold the information required in a relatively straightforward file compared to audio formats, so either of the these formats can be used. Additionally, both Kern and MusicXML have the advantage over audio formats in that less user input is required to convert printed scores into notational formats. However, Kern has a slight advantage over MusicXML due to the large database of Kern scores available, although the majority of Kern scores are also available in MusicXML. However, there are other websites that hold a small number of additional MusicXML scores such as Project Gutenberg (Project Gutenberg, n.d.). Additionally, both *Finale* and *Sibelius* can convert MusicXML documents into music notation. Therefore, either Kern or MusicXML would be appropriate for the present research.

3.5 Computer-Aided Music Analysis Technologies

The final factor that needs addressing when searching for patterns in music using computers is which technology to use (which is partly dictated by encoding formats). There are a number of different computer programs that have been developed to help aid music analysis that range from specific applications, such as Tonalities developed by Pople (2002) that analyses harmonic structures and tonal centres in music, to generalised analytical tools such as the Humdrum Toolkit developed by Huron (1997). The present research requires a technology that is able to cope with large datasets, which can search for patterns using only a general concept of a pattern, and which can produce statistical information on the patterns found. Currently, there are two main systems that do such generalised pattern searching and that might be appropriate for the present research: Melodic Match and the Humdrum Toolkit.

Melodic Match was developed at the University of Melbourne, with the first official release in 2007 (Wheatland, n.d.). Its purpose is to allow for the search of patterns in music notation and lyrics using a graphical interface based on MusicXML documents, and it is principally aimed at music researchers and analysts. The various different searches can either produce statistical information or provide location information for the patterns. The program has a number of different search features that facilitate searching for both melodic and rhythmic patterns with a number of different predetermined transformations, with or without criteria such as articulation. There is also the facility to search for a number of different patterns at the same time. The results are set out either as statistical information or as a coloured graphical representation of the score.

The advantages to this system are that it can search for both absolute pitch and relative intervallic values, and can search through multiple files as part of a single search. There are, however, a number of limitations in respect to the present research. Firstly, it can only match on exact patterns; i.e., it has no facility to include 'wild cards', and is therefore unable to match similarity between variant patterns beyond the set transformations. Secondly, it can only deal with absolute values for rhythms; i.e., it will regard four crotchets in 4/4 time as being different from four quavers in 4/8 time. Thirdly, it deals with MusicXML documents as separate files; searching through a large number of files simultaneously is very time-consuming and processor-intensive.

The Humdrum Toolkit was, like Melodic Match, designed as an aid for musicologists to help analyse music. The package consists of two components. Firstly, there is a predefined syntax using ASCII text that all encodings must follow. This specifies how the encodings should treat different properties of the music and its metadata (Huron, 1994, p. 8). A number of different translators have been written to convert different encoding formats such as MIDI into a suitable format for use by Humdrum Toolkit (Huron, 2002, p. 11). Secondly, there is the Humdrum Toolkit itself which has over seventy separate commands for the manipulation of data. These commands are small programs that perform a single operation to manipulate data, which can be welded together to create more complex manipulations (Huron, 2002, p. 14).

The main advantage of the Humdrum Toolkit over Melodic Match is that because there is an array of different commands available, there is the ability to fuse commands together, which makes it a very powerful tool for music analysis. However, this complexity does come at a cost, with Huron admitting that learning the commands can be like learning a computer programming language. Huron also believes that 'programmers will readily understand that the command-oriented structure is Humdrum Toolkit’s principal strength, because it allows complex scripts to be written and embedded in one’s favorite programming language' (2002, p. 21). Some researchers have tried to reduce the complexity of the commands by creating Graphical User Interfaces, such as Kornstädt's JRing and Taylor's Humgui (Huron, 2002, p. 25). However, neither of these systems offers an interface for all of the commands available under the Humdrum Toolkit.

Using the Humdrum Toolkit to tackle the present research project would have theoretically been possible, but it would have created a number of problems. Firstly, this project is not what the Humdrum Toolkit was created for. Although it was designed to allow both general and specific applications, with questions relating to music such as 'Which of the Brandenburg Concertos contain the B-A-C-H motive?' or 'In the music of Stravinsky, are dissonances more common in strong metric positions than in weak metric positions?' (Huron, 1998, p. 1), it was not specifically designed for mass data analysis on the scale of the present research. Secondly, the Kern documents are all semi-structured text-based documents without any indexing, making them difficult to query and computationally inefficient. Thirdly, a number of documents used for the present research were either in pdf format or were scanned from a printed score, which would have created difficulties in translating them into the Kern format (this would have involved converting them into MusicXML and then converting MusicXML into the Kern format).

As with using the Humdrum Toolkit, it would have theoretically been possible to conduct the present research using the MusicXML documents directly. There are a number of tools that allow the querying of XML documents such as XQuery or XPath, both of which were designed to query collections of XML data (Prigmore, 2008, pp. 126-127). However, XML is really a data-interchange format (i.e., it allows the porting of data from one program to another program without any further intervening conversions to the data or its structure), rather than a fully functional data-manipulation system. XML also has the same problem as Kern documents, in that all the documents are text-based making them computationally inefficient to query.

An alternative solution to using Melodic Match, the Humdrum Toolkit, or directly querying MusicXML documents, is to use a technology that was designed for mass data storage and manipulation, regardless of whether the technology was specifically designed for music data. The present research must search through large amounts of data to find patterns in the pitches and rhythms of music scores that are encoded using letters and numbers. Therefore, providing the technology used can search for patterns in letters and numbers within large datasets and produce suitable statistics, then it does not need to be specifically related to music.

Relational Databases are an obvious technology to choose, because they can hold large amounts of data which can then be queried and manipulated using Structured Query Language (SQL). According to Prigmore, a relational database is 'a persistent, self-describing, structured collection of related items of data' (2008, p. 10). The important concepts in Prigmore's definition are that the database is *self-describing,* because this means that it holds details of the structure of the information; it is *structured* meaning that all the information is stored in a consistent and organised way, i.e., every piece of data has its own place within the structure of the database, and it holds *related items of data,* meaning that all the pieces of information have some connection.

The advantages for the present research in using relational databases concern the speed of information retrieval, the use of a standard data-query language, and the concept of data mining. Relational databases have the advantage over the Humdrum Toolkit with speed of searching because they can implement indexing. Indexing data is much like indexing a book, and when implemented effectively can make searching for the required information quicker than non-indexed data (Prigmore, 2008, p. 510). Querying information in relational databases is also less complex than in the Humdrum Toolkit because databases use an internationally recognised language for querying data, the Structured Query Language (SQL), initially developed at IBM in the 1970s by Donald Chamberlin and Raymond Boyce, which was adopted by the International Organisation for Standards (ISO). SQL can be used to manipulate both the data and the meta-data within the database. This language has been implemented across a number of different Database Management Systems (DBMS), including *Oracle* and *SQLServer*, with Elmasri and Navathe arguing that it is one of the reasons for the success of the relational database model (2007, p. 233). Finally, relational databases also have the advantage of being able to implement data mining techniques easily, which, according to Elmasri and Navathe, are powerful tools for analysing large datasets that look for patterns and relationships (2007, p. 25).

The disadvantage with the relational database model for the present research is that the music is not in a suitable format. It is possible to transfer the pdfs and scanned scores into a tabular format straight away, however this would have involved writing a program for the Optical Character Recognition (OCR) of music that could then transfer the data directly to a relational database because one is not currently available. Of the pre-existing encoding systems, both Kern and MusicXML documents can be translated into tabular form for the relational database model. Here, MusicXML has the advantage over Kern in that algorithms already exist for importing XML documents into relational databases. Additionally, programs exist for the OCR of pdfs and scanned documents that import the data into MusicXML, unlike for Kern. Using these translation tools together with existing MusicXML documents means only one program would be required for the transfer of the data into the relational database. Therefore, due to the ease of importing MusicXML documents into a relational database, as well as there being programs already available to read and translate pdfs and scanned scores into MusicXML, the MusicXML was used for the present research.

3.6 The Four Factors

The requirements for a system to analyse patterns in music has been formulated using factors put forward by Uitdenbogerd and Zobel (2004). Ultimately, four factors have been identified as being necessary for the present research: defining the properties of patterns, defining similarity between patterns, sourcing or creating a suitable electronic encoding system for music, and sourcing or creating suitable computer programs.

A working definition of a pattern has been arrived at in Section 2.5.3 above, which provides a suitable definition that can be translated into code that a computer can process. Additionally, a definition of similarity has been formulated in Section 3.3 above based on selected properties of patterns that can also be translated into code. MusicXML was chosen as a suitable encoding system due to it providing all the information required, having a number of sources already encoded, being readily integrated into music editing software such as *Sibelius* (as discussed in Section 3.4), and being easily integrated with a relational database. Finally, a relational database model was chosen to implement the search algorithms due to its data-handling abilities, its use of SQL, and its ability to implement data mining techniques.

Consequently, the definitions for a pattern in music and for pattern similarity in the present research will be implemented using data mining techniques on MusicXML documents that have been uploaded into a relational database.

3.7 Summary

At the beginning of this chapter, it was shown that a variety of encoding systems for music and systems for analysing music have been developed. The majority of these encoding and music analysis systems are for specific tasks (such as analysing Bach Chorales), but there are some systems that are more generalised in their concept, such as Kern, MusicXML, the Humdrum Toolkit, and Melodic Match. However, the question was whether these systems were suitable for the present research.

It was argued that for the present research four factors, based on Uitdenbogerd and Zobel's six factors for using computers to match patterns (2004, p. 1054), were required. Firstly, a definition of a pattern that could be implemented using computers was required. In Section 2.5.3, a working definition of a pattern in music was provided for the present research, which can be translated into an algorithm for a computer to process. Secondly, a definition of similarity between two patterns was required. A definition of similarity is extremely complex, in that it covers a number of diverse approaches, from cognitive to analytic, and from motivic to structural, etc. There are a number of different approaches to implementing searches for similar patterns, such as weighting systems, neural networks, and earth-mover's distance algorithms. However, the investigation, implementation and testing of these algorithms for the task required is complex and time-consuming, and consequently beyond the scope of this research. Therefore, an alternative approach to matching patterns by comparing their length, their overall shape, their pitch/duration centres, and the distance between the highest and lowest or longest and shortest notes was devised (see Section 3.3 above).

The next factor was determining a suitable encoding system. A number of encoding systems were discussed, from audio- to notation-based systems. However, the requirement for the present research was that the encoding system should be easily able to handle pattern matching, which ruled out audio-based systems due to the problems in decoding the frequencies and instrumentation; and should have a large number of existing encodings already available, which ruled out SMDL. The remaining two notation-based encoding systems (Kern and MusicXML) can both handle pattern matching in terms of the present research, as well as having a selection of documents already encoded. Ultimately, the decision to use MusicXML was based on the fact that it is a data-interchange format, meaning that programs such as *Photoscore* and *Sibelius* are able to handle it, that there are a number of compositions already encoded in MusicXML, and that there are already algorithms for uploading the documents into a relational database.

The final factor that required addressing was which computer analysis system to use. The two obvious choices of the Humdrum Toolkit and Melodic Match were discounted. Melodic Match was discounted due to its inability to handle the definition of a pattern developed here, and the Humdrum Toolkit due to the complex nature of the system and the amount of processing power it would have required to mine the data. An alternative solution of using a relational database model was proposed. The advantages of the relational database model are that it can handle large amounts of data, it has an ISO standard querying language, and due to its indexing capabilities it is less processor intensive than Melodic Match and the Humdrum Toolkit.

4 Chapter 4: Using Knowledge Discovery in Databases

4.1 Introduction

Data mining is just one phase of the process of Knowledge Discovery in Databases (KDD). Elmasri and Navathe describe a total of six phases for KDD: Data Selection (deciding which data to use for the process), Data Cleansing (removing any data that is incorrect or corrupt), Data Enrichment (adding information that is not in the original data that will help with the understanding of the data), Data Transformation (using encoding to reduce the amount of data), Data Mining (processing and analysing the data), and Data Reporting (retrieving the results) (2007, pp. 946-947). There is, however, no consensus for the number of phases for KDD. For example, Maimon and Rokach have nine phases for KDD: Understanding the domain, selecting and creating the dataset, pre-processing and cleansing the data, transforming the data, selection of data mining technique, selection of an appropriate algorithm, running the algorithm, evaluating the results, and putting into practice the knowledge discovered (2005). However, both sets of phases can be reduced to three overall stages: data preparation, data mining, and data evaluation. The present research will implement these three stages overall, using Elmasri and Navathe's six phases. These phases will be followed in the present research because they all have an impact on, and are integral to the success of, the data mining process at the centre of its methodology.

This chapter gives an overview of how KDD works and how it is applied in this study. Each stage of the KDD process is discussed, together with its implementation in terms of the present research. Additionally, a section on how the system was tested is included. This explains how the code was tested to ensure the results were accurate by manually testing small samples of data. Also, a pseudo-random sample of data was created to test the methodology, together with three pre-determined patterns. The results from the pseudo-random and seeded pseudo-random testing show that any interference resulting from the methodology is negligible.

4.2 Data Selection

According to Temperley, 'Music notation provides a representation which is convenient for study and can also easily be converted into a format suitable for computer analysis' (2001, pp. 2-3). Unfortunately this statement does not really do justice to the difficulties involved in converting scores into an electronic format. Programs are available to assist in the conversion (e.g., *Photoscore* and *SharpEye*), however they are not totally accurate and require varying amounts of intervention in the conversion process depending on issues such as the quality of the scan and the complexity of the music. To minimise the use of conversion software, the repertoire chosen for the present research was largely determined by the availability of scores already in electronic format.

There are both musical and technical issues to consider when selecting the repertoire. Firstly (as already noted), the availability of scores in electronic format needs to be investigated in order to minimise the use of conversion software. Secondly, the repertoire needs to extend across different time periods in order to test whether patterns exhibit the evolutionary processes as well as the replicator properties. Testing whether these processes and properties exist can only be achieved by looking at the progress of patterns over an extended period of time. Thirdly, the repertoire should consist of only one genre. This is important because instruments and genres have their own idiosyncrasies (e.g., the pitch range of the instruments involved, etc.) that composers take into account, thereby possibly mediating the interchange of patterns between compositions. Genres also have their own individual styles and structures based on historical origins and development. Consequently, each genre could have its own set of genre-specific memes.

There is a large selection of scores in MusicXML format on the Kern Score website (Center for Computer Assisted Research in the Humanities, n.d.). The website lists fifty-one composers with a total of 108,703 files between them in an electronic format. However, the range of scores available is variable; for example, there are many encodings of Bach chorales but only seven symphonies (all by Haydn). Many of the scores are in multiple formats, including Kern and MusicXML. Other websites contain MusicXML documents, such as Project Gutenberg (n.d.), but such websites have a narrow range of repertoire in this format. Scores in pdf format are also available on the Project Gutenberg site, and on other sites such as the Petrucci Music Library (International Music Score Library Project, n.d.), which contains an extensive range of repertoire. Because none of these sites individually will provide enough suitable repertoire in MusicXML format for meaningful results, all will be drawn on, together with some conversions from pdf format and some scanned scores to fill gaps in the repertoire that is not available electronically.

Because the Kern Score website contains the most scores, this site helped determine the choice of repertoire. The site provides a number of different genres, including string quartets and piano sonatas. Because the chosen genre needed to encompass different time-periods, string quartets were selected to provide a representative example; composers from the classical period to the present day have written in this genre. String quartets also have the advantage of having mostly a single melodic line per instrument, making the pattern generation less complex. However, the number of string quartet scores available in Kern format was limited to a few composers (just five), and there are no comprehensive collections. Therefore, an additional source, such as the Petrucci Music Library, was required, upon which the necessary conversion process could take place in order to provide a representative sample of the genre across different time periods. Additionally, some scores were scanned to supplement the range of repertoire used.

This part of the research began by selecting MusicXML scores available from the Kern Score website. These covered 5 composers and a total of 289 movements from 81 compositions (see Table 4.1).

Composer	Number of Compositions	Number of Movements
Haydn	40	139
Mozart	23	73
Beethoven	16	71
Schubert	1	4
Brahms	1	2

Table 4.1: Number of compositions and movements downloaded from the Kern Score website

There were more documents than those listed above available on the Kern Score website, but there were problems with the format of some of these. For example, the file for Haydn's String Quartet op. 20 no. 6, 1st movement, had the first violin part split across two staves (see Example 4.1, where XPart 3 and XPart 4 both represent the 1st violin, XPart 0 is for the cello, XPart 1 is for the viola, and XPart 2 is for the 2nd violin part).

Example 4.1: Notation based on the MusicXML document of Haydn's op. 20 no. 6 1st movement bb. 1-4, downloaded from the Kern Score website

Some of the MusicXML files downloaded from the Kern Score website needed to be amended in order for *Sibelius* to be able to read them (as part of the data cleansing phase,) because *Sibelius* was unable to recognise the Umlauts in the documents. For example, most of the Mozart scores used an Umlaut in 'Härtel' in the 'YOR' key section (lines of code that provide metadata) of the document which states the original published score from which the encoding was made, i.e.

<!-- INFO key="YOR" value="Breitkopf & Härtel Edition" -->

This was easily rectified by removing these 'keys' from the affected documents because they are for information only and are not essential for the reading of the documents.

Time was also a limiting factor in how many of the available documents on the Kern Score website were used. Although the documents from the Kern Score website were already in the MusicXML format, a certain amount of preparation and checking of the documents was still required. Because time was limited, it was decided to restrict the number of documents used from the Kern Score website in order to concentrate on getting a greater spread of composers from other sources.

A small number of additional MusicXML documents were downloaded from the Project Gutenberg website. These covered just 2 composers, with a total of 5 movements from 2 works (see Table 4.2).

Composer	Number of Compositions	Number of Movements
Brahms	1	1
Schumann	1	4

Table 4.2: Number of compositions and movements downloaded from the Project Gutenberg website

Another source of documents was the Petrucci Music Library. Documents downloaded from this website were in pdf format and covered 15 composers with a total of 117 movements from 54 works (see Table 4.3).

Composer	Number of Compositions	Number of Movements
Haydn	19	24
Mozart	12	13
Beethoven	1	1
Schubert	3	12
Mendelssohn	3	12
Schumann	2	8
Smetana	1	2
Borodin	2	8
Brahms	3	9
Tchaikovsky	2	8
Grieg	1	4
Debussy	1	1
Sibelius	1	5
Ravel	1	4
Bartok	2	6

Table 4.3: Number of compositions and movements downloaded from the Petrucci Music Library website

Because the documents from the Petrucci Music Library were all in pdf format, they needed to be converted into MusicXML. This conversion was done using Neuratron's *Photoscore Ultimate* software. The conversion accuracy from the pdf document into *Photoscore* was rather disappointing, in part because of the poor quality of the pdf documents, and resulted in a large amount of manual correction being required (see Figure 4.1).

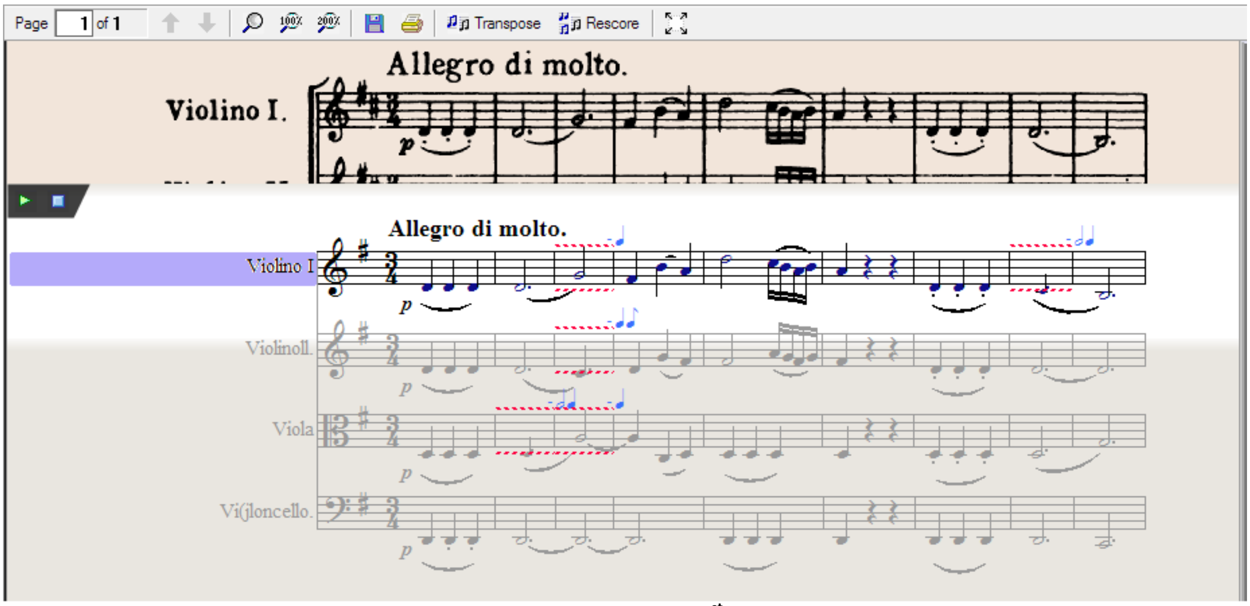

Figure 4.1: ***Photoscore*** **rendering of Haydn String Quartet op. 20 no. 4 1st movement bb. 1-9**

Another problem with using *Photoscore* was that there were errors in the conversion from pdf via *Photoscore* directly to MusicXML, specifically in some octave placements and incorrect accidentals (see Figure 4.2). As a result, *Photoscore* could not be relied upon to make the conversion to MusicXML. Fortunately, *Sibelius* can read the *Photoscore* proprietary files which could then be checked, and by using Recordare's *Dolet* plugin, can then convert the files into MusicXML. No issues concerning the conversion of pdfs via *Photoscore* then via *Sibelius* before translation to MusicXML via *Dolet* were encountered.

Figure 4.2: Incorrect D# in bars 338 and 339 during direct conversion from *Photoscore* to MusicXML: Bartok String Quartet no. 1 2nd movement bb. 337-346

Unfortunately using the Petrucci Music Library still only provided documents for a limited number of compositions. Therefore, additional works in printed-score form were scanned into *Photoscore,* which were then converted by *Sibelius* via the *Dolet* plugin into MusicXML. This process covered 6 composers with a total of 31 movements from 9 works (see Table 4.4).

Composer	Number of Compositions	Number of Movements
Debussy	1	3
Dvorak	2	8
Janacek	1	4
Prokofiev	1	3
Shostakovich	3	11
Smetana	1	2

Table 4.4: Number of compositions and movements scanned into *Photoscore*

The same problems that *Photoscore* had with pdf documents were also encountered with the scanned printed documents. The quality of the conversion into *Photoscore* was sometimes poor and *Photoscore* would again not correctly convert to MusicXML. As a result, additional manual intervention was required within *Photoscore*, together with using *Sibelius* and *Dolet* to convert the documents to MusicXML.

In total, documents were either downloaded or scanned for 19 composers with a total of 442 movements from 111 compositions (see Table 4.5). The earliest composition used was written in 1762 and the last was written in 1970, allowing for a wide time-frame for the evolutionary processes of selection, replication and variation to act. A complete list of the compositions used and their sources is given in Appendix 1 - Working Definitions.

This selection of works and movements is by no means a comprehensive representation of the genre as a whole. A number of composers who wrote string quartets are missing, such as Boccherini, Britten, Schoenberg, Spohr and Villa-Lobos, all of whose scores are not easily available electronically. It is also rather biased towards composers such as Haydn, Mozart and Beethoven, who have a large selection of their output on the Kern Score website. Additionally, only Beethoven and Mozart have had all of their string quartets transferred to the database used for the present research. The selection of composers and works used was partly due to the availability of the material, partly due to the problems of using *Photoscore*, and partly due to the time-consuming nature of scanning and converting scores into MusicXML and the resultant time pressures on the research. However, the sample does cover a range of composers across a wide time-span that will provide meaningful and interesting results that can then be explored within the remit of the present research.

Composer	Number of Compositions	Number of Movements
Haydn	40	163
Mozart	23	86
Beethoven	17	72
Schubert	4	16
Mendelssohn	3	12
Schumann	3	12
Smetana	1	4
Borodin	2	8
Brahms	3	12
Tchaikovsky	2	8
Dvorak	2	8
Grieg	1	4
Janácek	1	4
Debussy	1	4
Sibelius	1	5
Ravel	1	4
Bartok	2	6
Prokofiev	1	3
Shostakovich	3	11

Table 4.5: Total number of compositions and movements used for the dataset

4.3 Data Cleansing

The next phase of KDD is data cleansing, which involves ensuring that all the data used is accurate. In this case, all the electronic documents needed to be checked to ensure that both the encoding had been done correctly, and that the format of the documents was correct. This was achieved by checking that the MusicXML documents could be read by *Sibelius*, and then by checking the resultant *Sibelius* score against a printed edition.

In the previous section on data selection, it was noted that *Photoscore* caused problems in its reading of pdf documents and its conversion into MusicXML. Unfortunately these were not the only problems encountered whilst using this software. There were also some areas where it was unable to cope with certain aspects of the music.

Photoscore was unable to read or encode tremolo markings, such as in Brahms' String Quartet op. 51 no. 1, 1st movement, bb. 1-6 in the viola and cello parts (see Figure 4.3). To correct this would either have involved altering the notes to the correct value manually without the tremolo marking in *Photoscore*, then opening the *Photoscore* document in *Sibelius* and adding in the tremolo marking, or expanding out the tremolo markings that had specific rhythmic values. This would have been extremely time-consuming because constant reference back to the score would be required to ensure that all the tremolos were correctly encoded. It was therefore decided to remove the tremolo marking and leave them as single notes with the appropriate rhythmic values (see Example 4.2). For the sake of consistency, tremolo markings were also removed from all documents that were not generated by means of *Photoscore*. This will have an impact on the pattern generation, for the number of repeated same-note patterns identified by the system will be fewer as a result of not including the tremolos.

Figure 4.3: *Photoscore* reading of the tremolo in Brahms String Quartet op. 51 no. 1 1st movement bb. 1-6

Example 4.2: Interpretation of the tremolo in Brahms String Quartet op. 51 no. 1 1st movement bb. 1-6

Difficulties also arose over how to handle repeats in music using the *Photoscore* software. *Photoscore* has the ability to put in repeat marks but not first or second time indicators (e.g., Beethoven String Quartet op. 135, 2nd movement, b. 67), or *Da Capo* repeats (e.g., at the end of the trio section of the minuet and trio movement from Haydn's String Quartet op. 42). *Photoscore's* inability to handle certain types of repeats called into question whether the pattern generation system should take repeats into account. Repeats tend to be structural and because structural memes are beyond the remit of the present research they do not need to be included. However, some composers write out repeats or *Da Capos* if they want to make subtle changes (as in the second movement of Beethoven's String Quartet op. 135 where the *Da Capo* is written out), in which case repeats should be included. Therefore, some of the scores will have the repeats marked whereas others will have the repeats written out.

Repetition is important for the propagation of memes so not including repeats could have an impact on showing possible memes. However, repeats are not always observed in performance and sometimes propagation takes place through performance rather than scores. Unfortunately, the overriding factor for the present research was, as with the tremolos, the time-consuming nature of restoring the repeats that *Photoscore* could not handle, and therefore repeats were also left out. For the sake of consistency, repeat markings were removed from those documents not generated through *Photoscore*. It is recognised that this still leaves the problem of some composers writing out the repeats, which have been left unaltered. Although this is not ideal, it was necessary to keep the research focused.

Ornamentation was also a problem for *Photoscore*. Symbolic ornaments such as trills, mordents etc. were not easily picked up by *Photoscore*, with many being interpreted as plain text. The program also had difficulty in associating the ornaments with the correct notes, with some being linked to the incorrect stave. Grace notes also caused problems in that they would be read as normal notes, thus creating an incorrect number of beats in a bar. Acciaccaturas were a particular problem, in that *Photoscore* would interpret the slash as part of the stem, thus altering the length of the note. However, the definition of a meme in music for the purposes of the present research ignores all symbolic ornamentation (see Section 2.5.3 above). As such, most of the ornamentation was either ignored where it did not affect the notes (such as trill markings), or removed where it did affect the notes (such as acciaccaturas).

The next problem encountered in using *Photoscore* was how the software handled tied notes and slurs. Sometimes it would confuse tied notes and slurs; e.g., two identical consecutive notes with a tie over them in the original would be encoded as two identical consecutive notes with a slur rather than a tie (see Figure 4.4). This was a problem for the present research because it was important to distinguish whether a note was sounded more than once when generating the possible patterns. Unfortunately, the *Photoscore* interface does not make it clear whether it has encoded the marking as a slur or a tie, which meant checking each marking that covered two identical consecutive notes to determine if it was encoded as a tie and amending where appropriate.

Figure 4.4: ***Photoscore*** **rendering of ties/slurs in bars 28 and 29 of the cello part in Brahms String Quartet op. 51 no. 1 2nd movement bb. 27-31**

Using *Photoscore* to convert pdfs or scanned scores into MusicXML via *Sibelius* and *Dolet* also raised issues concerning the quality of the editions of the music in pdf. Although the present research is concentrating on pitches and rhythms and ignoring secondary parameters, there were still the occasional problems encountered from editorial errors in editions from which the pdfs and scanned scores were taken. These covered incorrect number of beats in the bar (e.g., Haydn's String Quartet op. 64 no. 2, 1st movement, b. 20, where the violin 2 part has too many quaver rests), to missing clefs (e.g., Mozart's String Quartet K. 575, 4th movement, b. 8, where the cello part is missing the return to bass clef after beginning the movement in the treble clef). Where these mistakes were detected, they were corrected via the *Photoscore* software.

It was not just editorial errors that were detected. A number of encoding errors were found in the documents downloaded from the Kern Score website. Some of these errors were a result of mistakes in the conversion from the original Kern document to the MusicXML document, and others were incorrect in the original Kern document and therefore also in the MusicXML document.

Examples where there were errors in the MusicXML documents but not in the original Kern document included problems with accidentals (e.g., Beethoven's String Quartet op. 130, 1st movement, b. 56, where the violin 2 part has a B natural that should be a B double-flat), problems with key signatures (e.g., Beethoven's String Quartet op. 130, 1st movement, b. 72, where the key signature should have six flats and not five as in the MusicXML document), instruments in the wrong order (e.g., Haydn's String Quartet op. 54 no. 3, 1st movement, were the instruments are in the reverse order with the cello part at the top), problems with tuplets (e.g., Brahms' String Quartet op. 51 no. 1, 3rd movement, b. 40, where the triplets are not encoded leaving an incorrect number of beats in the bar), and notes missing (e.g., Mozart's String Quartet K. 170, 3rd movement, bb. 5-6, where the violin 1 part has all the notes missing).

Examples where the errors appeared in both the MusicXML and Kern documents include missing accidentals (e.g., Beethoven's String Quartet op. 18 no. 1, 1st movement, b. 70, where the violin 2 part has an E that should be an E flat), incorrect notes (e.g., Beethoven's String Quartet op. 18 no. 1, 1st movement, b. 257, where the viola part has two B flats that should be As), missing appoggiaturas (e.g., Haydn's String Quartet op. 64 no. 4, 1st movement, bb. 19-20, where the violin 1 part has appoggiaturas missing from the C sharps), key signatures missing from trios (e.g., Mozart's String Quartet K. 171, 2nd movement), and incorrect enharmonic equivalents (e.g., Beethoven's String Quartet op. 18 no. 2, 4th movement, b. 343, where the cello part has an E sharp that should be an F natural).

Some of these errors were the result of the editions used having errors (e.g., Mozart's String Quartet K. 173, 3rd movement, b. 8, where the cello part has a C that should be a B flat to tie in with both the harmony of the bar and because it is part of a parallel octave passage with the viola part). However, most that were investigated were the result of either being incorrectly encoded in the Kern original, with the error therefore carried over to the MusicXML document, or being incorrectly converted to MusicXML.

Having found some issues with using *Photoscore* and pre-existing encodings, all of the encodings were examined within *Sibelius* and, if necessary, re-encoded via the *Dolet* plugin to ensure consistency of the MusicXML encoding and to correct any errors that were identified. Although a number of errors were discovered during this process, it is not certain that all errors have been rectified.

Accuracy will always be a problem when encoding music electronically, in the same way that it is a problem during the editing and creation of paper scores. Even when the scores are checked by specialist editors and proof-readers, there will always be occasional mistakes especially, when there are time-constraints. The latter was a significant factor in searching for errors within the encodings in the present research.

4.4 Database Design

A diversion from the KDD process is now necessary before a description of the next stage (data enrichment) because it is important to understand the structure of the relational database behind the KDD process together with how the information has been imported from the MusicXML documents into the relational database.

Relational databases are controlled by Database Management Systems (DBMS) that are usually proprietary. There are a number of different DBMSs available, with IBM, Oracle and Microsoft all producing their own, together with a number of open-source products such as MySQL. There are no overriding factors that need to be considered when choosing a DBMS for the present research because the database is not business-critical, nor will it be compiled for distribution. Therefore, the research uses Microsoft's *SQLServer* 2008 R2 Express Edition because this is available as a free download, has plenty of support from both Microsoft and a community of users, and I have knowledge of this technology, having used it to build a number of business administration systems.

The concept behind the relational database is that the data is split into a number of different tables that are linked together by 'keys'. These keys are unique identifiers that create different relationships between the tables. This means that when one table is investigated, it is easy to see the data in all the other related tables linked by the keys.

The present research uses tables that can be grouped into two distinct types. Firstly, there are the main tables that hold the data extracted from the MusicXML documents. Secondly, there are the tables for the data transformation and data mining processes. There is, nevertheless, an additional table (tblFilePath) that is used solely by the conversion program to ensure that the same file is not converted more than once.

4.4.1 Main Tables

The present research uses a central table called tblNote that holds details of all the pitches and their durations from the source repertoire: every note in every piece has a single row of data associated with it in this table (see Figure 4.5). The columns used in this table can be divided into four types: firstly the key that links the table to the other tables, secondly those holding details about the position of the notes within the scores, thirdly those holding details about the notes themselves, and fourthly those that will be filled by calculated values as part of the KDD processes.

	NoteID	PieceID	Movement	PartID	Bar	NoteOrder	Pitch	Octave	Duration	Tie	Voice	NewPart	Chord	RelValue	AbsValue	TimeValue
1	490541	28	4	P4	697	1	G	2	1536	N	1	0	2	S0	S0G	1
2	490542	28	4	P4	697	1	D	3	1536	N	1	0	3	NULL	NULL	NULL
3	490543	28	4	P4	697	1	B	3	1536	N	1	0	3	NULL	NULL	NULL
4	490544	28	4	P4	697	2	R	3	1536	N	1	0	1	R0	00R	1
5	490547	28	4	P4	700	1	G	2	7680	E	1	0	1	S0	S0G	5
6	490548	28	4	P4	700	2	R	2	1536	N	1	0	1	R0	00R	0.2
7	490549	29	1	P1	0	1	B-1	3	384	N	1	0	1	U15	U1B-1	0.25
8	490550	29	1	P1	1	1	B-1	3	768	N	1	1	1	B-1	S0B-1	1
9	490551	29	1	P1	1	2	D	4	384	N	1	0	1	U4	U0D	0.5
10	490552	29	1	P1	1	3	D	4	768	N	1	0	1	S0	S0D	2

Figure 4.5: A sample of the data from tblNote that consists of the last notes of the cello part of Grieg's String Quartet op. 27 and the first notes of the violin 1 part of Haydn's String Quartet op. 1 no. 1

There are two columns of the first type: the NoteID, which gives each note a unique numerical identifier; and the PieceID, which links tblNote to tblPiece through a numeric value. The second type of column gathers together those columns that describe the position of the note within the score and includes the Movement column, which holds the number of the movement the note belongs to; the PartID column, which refers to which instrument the note belongs to (P1 for Violin 1, P2 for Violin 2, P3 for Viola, and P4 for Cello); the Bar column, which relates to the bar number the note is found in; the NoteOrder column, which gives the order in which the note appears within a bar; and the Voice column, which gives a numerical indication for which voice the note belongs to when there is more than one voice in a part. The third type of column indicates the details of the note and includes the Pitch column, which contains the pitch name (together with '+' for sharps and '-' for flats); the Octave column, which gives the octave of the pitch; the Duration column, which gives a numeric value for how long the pitch lasts; and the Tie column, which indicates if the note is part of a tie (with 'N' for not a tie, 'B' for the beginning of a tie, and 'E' for the end of a tie). The fourth type of column includes all the calculated values: the NewPart column, which states whether the note is the beginning of a new instrument, movement or piece; the Chord column, which indicates if the note is a part of a chord; the AbsValue column, which indicates whether the note has gone up or down from, or remained the same as the previous note together with its pitch; the RelValue column, which indicates the relative interval from the previous note; and the TimeValue column, which indicates the relative value of the duration compared to the previous note (see Table 4.6).

The PieceID in tblNote connects to tblPiece, which has details regarding the piece in which this particular note occurs, including the piece title, year and composer. It includes linked and calculated columns together with details of the pieces. The columns are: the PieceID, which links to tblNote; the ComposerID, which links to tblComposer; the PieceTitle, which gives the relevant catalogue number of the work; the PieceYear, which gives a year for the work (this is discussed in Section 4.5 below); and the PeriodID, which is a calculated column depending on the PieceYear column (see Table 4.7)

Column Type	Column Name	Column Values	Example Value
Link	NoteID	Unique identifier	465623
	PieceID	Number linked to tblPiece unique identifier	28 (indicating link to PieceID 28 in tblPiece)
Note Position	Movement	Number of the movement	1 (indicating the first movement)
	PartID	P1, P2, P3 or P4	P2 (indicating Violin 2)
	Bar	Bar number	62 (indicating bar 62)
	NoteOrder	Number indicating position within the bar	1 (indicating the first note in the bar)
	Voice	Number	1 (indicating that it is in the top voice)
Note Detail	Pitch	Pitch of the note including number of sharps(+) or flats(-)	F+1 (indicating an F sharp)
	Octave	Number	5 (indicating the 5^{th} octave
	Duration	Number	1536 (providing a comparable figure against other duration values)
	Tie	N or B or E	N (indicating there is no tie)
Calculated Value	NewPart	1 or 0	0 (indicating that it is not the start of an instrument, movement or piece)
	Chord	1 = single line note, 2 = top note of violin or viola part, or bottom note of cello part, or 3 = any other note	1 (indicating it is not part of a chord within the part)
	RelValue	U, D or S followed by number of semitones from last note	D1 (indicating that the note is 1 semitone lower than the previous note)
	AbsValue	U, D, or S followed by number of octaves difference from last note followed by pitch of this note	D0F+1 (indicating that it has gone to the F sharp less than an octave down since the previous note)
	TimeValue	Number indicating how many times longer or shorter the note is compare to the previous note	1 (indicating that the note is the same length as the previous note)

Table 4.6: Description of the columns present in tblNote

Column Type	Column Name	Column Values	Example Value
Link	PieceID	Unique identifier	28
	ComposerID	Number linked to tblComposer unique identifier	7 (indicating link to ComposerID 7 in tblComposer)
Piece Details	PieceTitle	Catalogue Number	27-0 (indicating op. 27)
	PieceYear	Year	1877 (indicating that the composition was begun in 1877)
Link and Calculated Value	PeriodID	Number linked to tblPeriod	4 (indicating period 4) and linking to PeriodID in tblPeriod

Table 4.7: Description of the columns present in tblPiece

The ComposerID in tblPiece connects to tblComposer, which holds the names of the composers (see Table 4.8).

Column Type	Column Name	Column Values	Example Value
Link	ComposerID	Unique identifier	7
Detail	Composer	Name of composer	Grieg

Table 4.8: Description of the columns present in tblComposer

These are the three main tables in the database and they are connected via the NoteID, PieceID and ComposerID columns as shown in the entity relationship diagram (see Figure 4.6). The relational nature of the database means that it is possible to trace a path from a particular note to a particular composer by moving through the tables using the PieceID to connect tblNote and tblPiece, and then the ComposerID to link tblPiece to tblComposer.

In addition to these three main tables, there are two lookup tables (i.e., tables that store repeated data which can be given a unique identifier): tblNoteLookup, which gives all enharmonic equivalents the same reference number; and tblPeriod, which gives the range of years that each period covers.

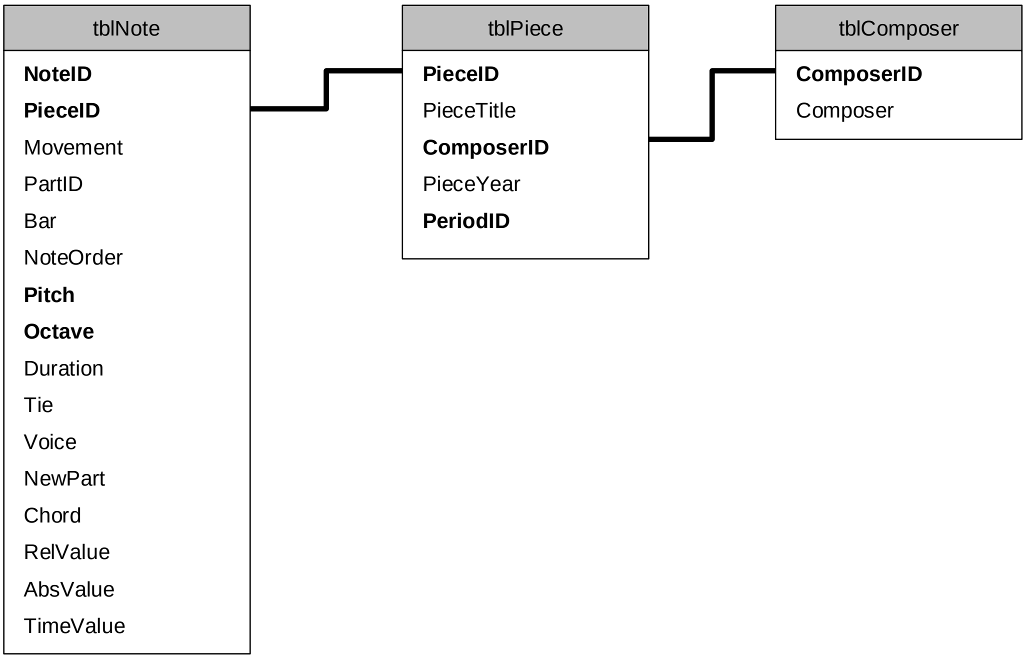

Figure 4.6: Entity relationship diagram with tblNote, tblPiece and tblComposer

In order to work out the relative intervallic distance between consecutive notes, the system needs to know when two notes with different pitch name values are enharmonic equivalents. This is achieved using tblNoteLookup, which lists all the different pitch values (Pitch column) together with their octaves (Octave column), giving all enharmonically equivalent pitches a unique number (NoteLookupID column) (see Table 4.9).

Column Name	Column Values	Example Value
NoteLookupID	Number	55
Pitch	Pitch of the note including number of sharps(+) or flats(-)	F+1 (indicating an F double sharp)
Octave	Number	5 (indicating the 5th octave)

Table 4.9: Description of the columns present in tblNoteLookup

Pieces are grouped together by periods according to the year in which the composition was started (see Section 4.6.1). These periods are controlled via tblPeriod, which lists the years covered by each period using PeriodStart and PeriodEnd columns (see Table 4.10).

Column Name	Column Values	Example Value
PeriodID	Unique Identifier	4
PeriodStart	Year	1851
PeriodEnd	Year	1903

Table 4.10: Description of the columns present in tblPeriod

These two tables can be fitted into the entity relationship diagram by linking tblNote to tblNoteLookup on Pitch and Octave, and by linking tblPeriod to tblPiece on PeriodID (see Figure 4.7).

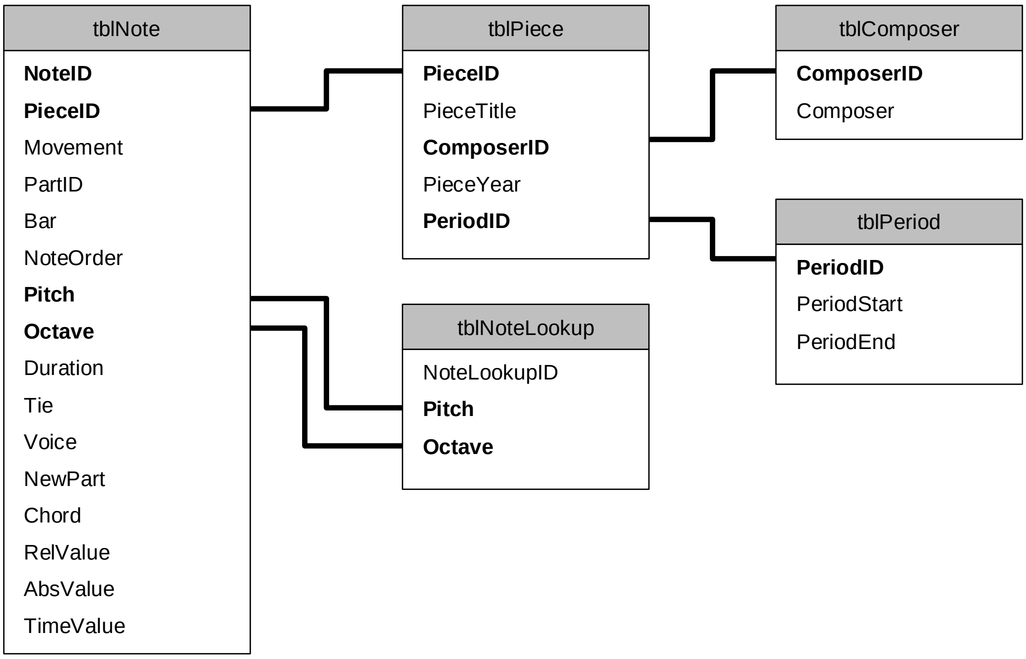

Figure 4.7: Entity relationship diagram with tblNote, tblPiece, tblComposer, tblPeriod, and tblNoteLookup

4.4.2 Data Transformation and Mining Tables

The data transformation and mining tables deal mainly with the pattern generation and ranking generation part of the database, with two tables for patterns (tblPatterns and tblPatternProperty) and one table for ranking positions (tblRanking). An additional table has been used to keep track of what stage the system is at (tblPass).

The main pattern table is tblPatterns that holds all of the three- to eleven-note patterns generated by the system (in accordance with the working definition of a pattern in music), using both the absolute pitch values and the relative intervallic values of the notes. It also links to tblPatternProperty using the AbsID, RelID and TimeID columns. Additionally, there are columns (beginning with DeDup) to show if the patterns should be included as part of the de-duplication process, as described in Section 5.3 (see Table 4.11).

Column Name	Column Values	Example Value
PatternID	Unique identifier	3620599
NoteID	Number	465623 (indicating that the pattern starts with NoteID 465623 from tblNote)
AbsValue01	Pitch of the first note in the pattern	F+1 (indicating the pattern starts with an F sharp)
RelValue01	0	Always a 0 (indicating the starting point for the relative values)
TimeValue01	1	Always a 1 (indicating the starting point for the relative values of duration)
AbsValue02…11	U, D, or S followed by number of octaves in change since last note followed by pitch of this note	D0F (indicating that it has gone down less than an octave to an F natural since the previous note)
RelValue02…11	U, D or S followed by number of semitones from last note	D1 (indicating that the note is 1 semitone lower than the previous note)
TimeValue02…11	Number indicating how many times longer or shorter the note is since the last note	1 (indicating that the note is the same length as the previous note)
NumNotes	Number	5 (indicating the pattern has 5 notes)
AbsID	Number	7042221 (indicating the PatternPropertyID in tblPatternProperty associated with this pattern)
RelID	Number	3033743 (indicating the PatternPropertyID in tblPatternProperty associated with this pattern)
TimeID	Number	647825 (indicating the PatternPropertyID in tblPatternProperty associated with this pattern)
DeDupA	Number	1 (indicating that this pattern should be included with the de-duplicated patterns – see Section 5.3)
DeDupAD	Number	1 (indicating that this pattern should be included with the de-duplicated patterns – see Section 5.3)
DeDupR	Number	0 (indicating that this pattern should *not* be included with the de-duplicated patterns – see Section 5.3)
DeDupRD	Number	0 (indicating that this pattern should *not* be included with the de-duplicated patterns – see Section 5.3)

Table 4.11: Description of the columns present in tblPatterns

The tblPatternProperty holds those properties of patterns that are used to determine pattern similarity (NumNotes, Shape, PitchCentre, and HighLow columns) and it links to tblPatterns on PatternPropertyID using AbsID, RelID and TimeID (see Table 4.12).

Column Name	Column Values	Example Value
PatternPropertyID	Unique identifier	7042221
PropertyType	A (for absolute pitch values), R (for relative intervallic values), D (for duration values)	A (indicating that the properties are for absolute pitch values)
NumNotes	Number	5 (indicating there are 5 notes in the pattern)
Pattern	Temporary coding for the pitches and/or durations of the notes in the pattern	Blank – (used as a temporary storage during the creation of pattern properties)
Shape	Combinations of U for up, D for down and S for same.	D (indicating the pattern goes downwards only)
PitchCentre	Number	-1 (indicating the average pitch movement)
HighLow	Number	4 (indicating there are 4 semitones between the highest and lowest notes)

Table 4.12: Description of the columns present in tblPatternProperty

Both tblPatterns and tblPatternProperty can be linked to tblNote as shown in Figure 4.8.

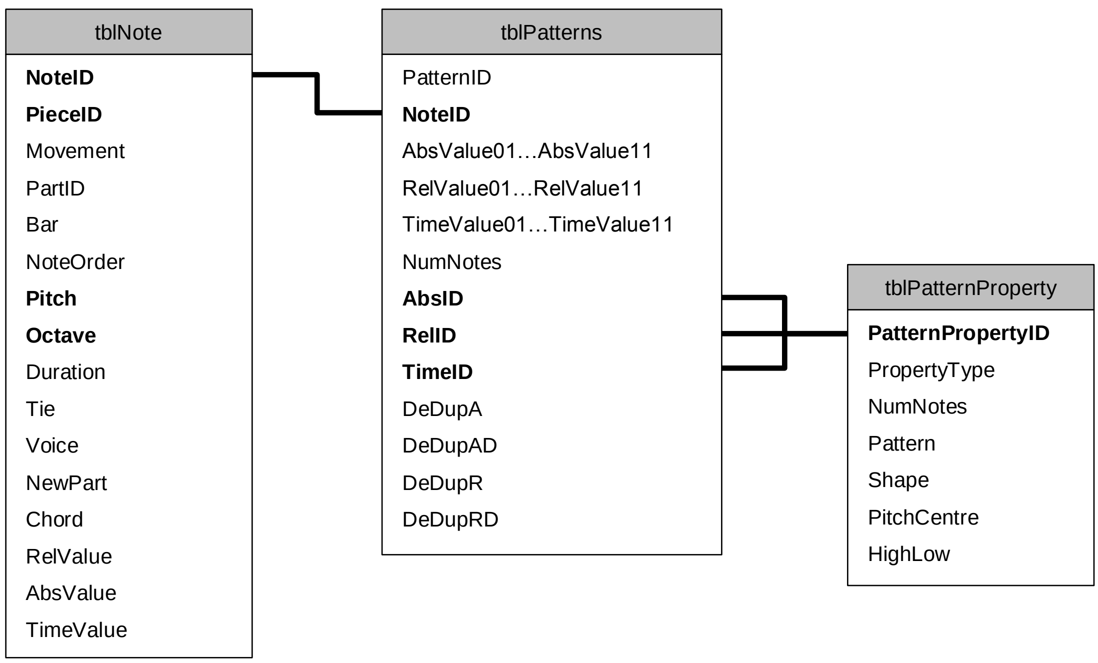

Figure 4.8: Entity relationship diagram with tblNote, tblPatterns, and tblPatternProperty

The next table, tblRanking, deals with ranking the patterns according to the frequency of their appearances within the dataset (see Section 4.7 below). There are a number of different ranking columns (all begin with Rank or NewRank), which relate to whether the ranking is across all pattern lengths, within each pattern length, or whether the ranking was calculated after the de-duplication process (see Section 5.3). Both the PitchID and TimeID columns refer to the PatternPropertyID in tblPatternProperty (see Table 4.13).

Column Name	Column Values	Example Value
PatternDescriptor	A (for absolute pitch values), AD (for absolute pitch values with duration), R (for relative intervallic values) or RD (for relative intervallic values with duration)	AD (indicating absolute pitch values with duration)
PeriodID	Number	4 (indicating period 4)
PitchID	Number	7042221 (indicating that it links to tblPatternProperty where PatternPropertyID = 7042221)
TimeID	Number	8575744 (indicating that it links to tblPatternProperty where PatternPropertyID = 8575744)
PatternType	Combination of pitches and durations	F+1:1/D0F:1/D0E:1/D0E-1:1/D0D:1 (indicating the pitches and durations of the notes in the pattern)
PitchShape	Combinations of U for up, D for down and S for same.	D (indicating the pitches of the pattern continuously progress downwards through the pitches)
DurationShape	Combinations of U for up, D for down and S for same.	S (indicating the durations of the notes are all the same)
NumNotes	Number	5 (indicating there are 5 notes in the pattern)
TotalNum	Number	26 (indicating that this pattern appears 26 times within this period)
Mvts	Number	7 (indicating that this pattern appears in 7 movements within this period)
RankAll	Number	963 (indicating that this pattern comes 963rd within the ranking positions for all pattern types)
RankNum	Number	117 (indicating that this pattern comes 117th within the ranking positions for pattern types of the same length)
NewTotalNum	Number	26 (indicating that this pattern appears 26 times within this period after de-duplication)

Column Name	Column Values	Example Value
NewMvts	Number	7 (indicating that this pattern appears in 7 movements within this period after de-duplication)
NewRankAll	Number	672 (indicating that this pattern comes 672nd within the ranking positions for all pattern types after de-duplication)
NewRankNum	Number	56 (indicating that this pattern comes 56th within the ranking positions for pattern types of the same length after de-duplication)
NewRankAllRelFreq	Number	906 (indicating that this pattern comes 906th within the ranking positions for all pattern types based on the relative frequency algorithm, explained in Section 5.5, after de-duplication)

Table 4.13: Description of the columns present in tblRanking

The tblRanking table can be linked back to tblNotes via tblPatterns tables as shown in Figure 4.9.

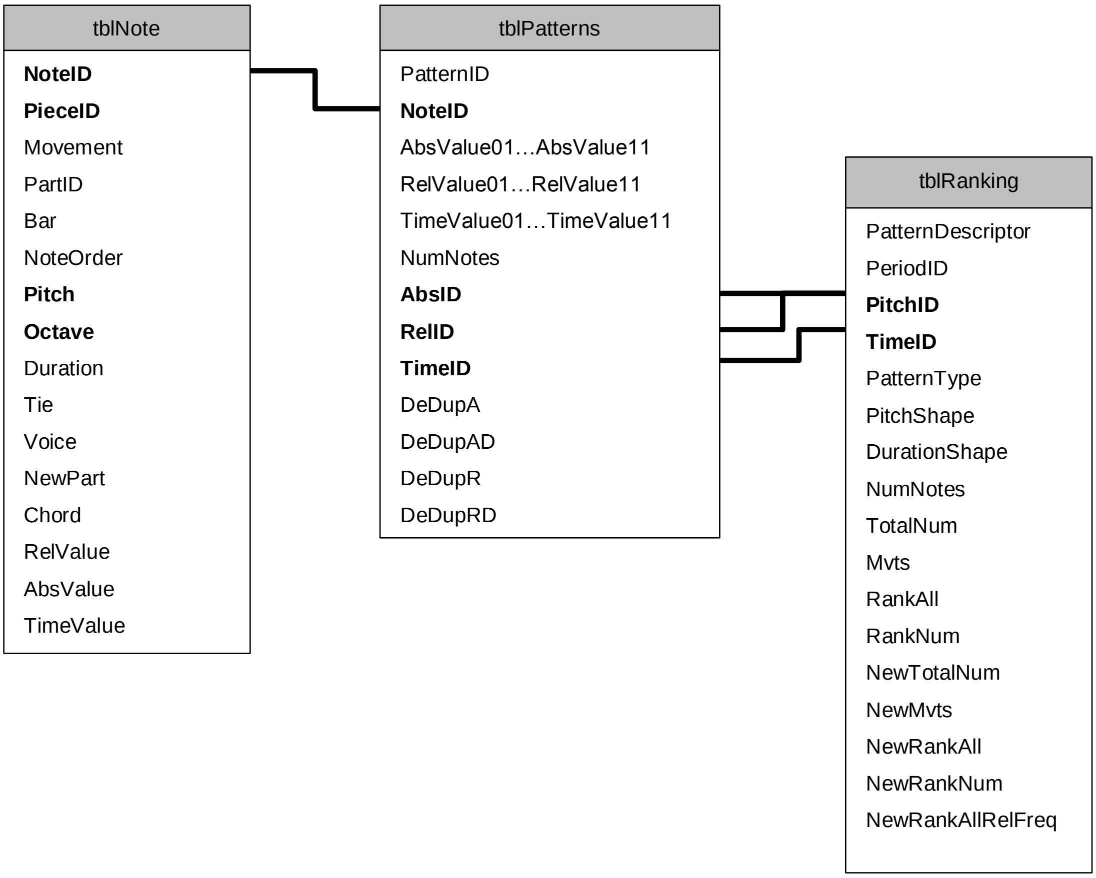

Figure 4.9: Relationship diagram with tblNote, tblPatterns, and tblRanking

The last table used by the pattern and ranking generation system is tblPass, which keeps a record of which pieces of code have been run (see Table 4.14).

Column Name	Column Values	Typical Value
Pass	Number between 0 and 6	1 (indicating that this record refers to the piece of code labelled Pass01)
Completed	Number between 0 and 6	1 (indicating that the associated stage of the code has been completed)

Table 4.14: Description of the columns present in tblPass

4.4.3 Additional Table

The final table in the database is tblFilePath, which holds a record of the file path of all the MusicXML documents that have been transferred into the Database (see Table 4.15). This is to ensure that the same file is not transferred more than once.

Column Name	Column Values	Example Value
FilePath	Path of an uploaded file	D:\Andrew\XML Files - NEW\Grieg\27-0-1.xml

Table 4.15: Description of the columns present in tblFilePath

A complete entity relationship diagram detailing all of the related tables can be found in Appendix 4.

4.4.4 Database Encodings

The database uses various codes to represent some of the data. These include codes for the pattern descriptors, the shapes of patterns, the pitches and durations of the patterns, and the pattern types.

Pattern descriptors are indicated using four different codes, representing different combinations of whether the pattern uses the absolute pitch or relative intervallic values of pitches, and whether they include duration as well as pitch (see Table 4.16).

Code	Meaning
A	The pattern uses the actual pitch values without any durations
AD	The pattern uses the actual pitch values together with relative durations
R	The pattern uses the relative intervallic values of pitches without any durations
RD	The pattern uses the relative intervallic values of pitches together with relative durations

Table 4.16: Pattern descriptor codes

The four pattern descriptors were chosen so that a comparison could be made between whether the absolute pitch of patterns produced different sets of results than relative intervallic value patterns, i.e., how important are the absolute pitch values to memes compared to relative intervallic values. Likewise, it was decided to separate out the pitch values from the rhythmic values to see if the pitch values without their durations produce a different set of results from pitch values with their durations, i.e., how important is the rhythmic element of patterns to memes. An alternative method would have been to use the scale degree classification of the pitches. However, this would have meant analysing the music for the different keys within each section of the pieces in order to convert the pitches to their appropriate scale degree classification, together with overcoming the major problems associated with producing an algorithm that can distinguish between the different musicological styles covered by the present research.

The shape directions of the pitches within patterns are indicated using three different codes representing whether the pattern is going upwards, staying on the same note, or going downwards. The same indicators are used to show if the durations of notes are getting longer, shorter, or are staying the same (see Table 4.17).

Code	Pitch Meaning	Duration Meaning
U	The pattern moves upwards in pitch	The duration gets longer
S	The pattern stays on the same pitch	The duration stays the same
D	The pattern moves downwards in pitch	The duration gets shorter

Table 4.17: Pattern shape codes

The absolute pitch values in tblNote are in the form of the note letter followed by a plus sign for sharp and a minus sign for flat. If there is just one flat or sharp the plus or minus sign is followed by the number 1; if it is a double sharp or double flat then the plus or minus sign is followed by the number 2. For example, A double sharp is indicated by A+2. For relative intervallic values, the number represents the number of semitones between the note and its predecessor. For example, the number 5 indicates that the note is five semitones either above or below the previous note.

Each note is preceded by a letter indicating the direction of the pattern from the previous note. U is for upward movement, D is for downward movement, and S means the note is the same as the previous note. In addition, for the absolute value of pitches a number will follow the direction indicating how many octaves difference there is. For example D1A means the pitch A is over one octave, but less than two octaves, below the previous note.

The pattern encoding in tblRanking is formed by placing all of the notes within a given pattern together in one field separated by forward slashes. Example 4.3 shows the encoding for a four-note A pattern descriptor pattern together with its musical notation.   The notes start on an F sharp, go up to an A, then go down more than 2 octaves to a G sharp, and ends by going up to the A next to the G sharp.

F+1/U0A/D2G+1/U0A = 

Example 4.3: Example of the encoding of a four-note A pattern descriptor pattern type

Likewise, Example 4.4 shows the encoding for a four-note R pattern descriptor pattern type together with its musical notation (starting on a G, although it can start on any note), where the note goes up 5 semitones, then up another semitone, then down 6 semitones to end on its starting note.

0/U5/U1/D6 = 

Example 4.4: Example of the encoding of a four-note R pattern descriptor pattern type

Durations can also be indicated in patterns using numbers to represent the duration relative to the previous note. Therefore the number 1 means the note is of the same duration as the previous note, values greater than 1 mean the note is longer than the previous note, and values less than 1 mean the note is shorter than the previous note. The duration is separated from the pitch by a colon, e.g., U6:2 means the note is 6 semitones above the previous note, and is twice as long as the previous note. Dealing with tuplets using relative values does not cause any problems. For example, a crotchet followed by a triplet quaver pattern would be encoded as 1 for the crotchet, because it is the first note in the pattern, then 0.3333, because the second note is a third of the value of the first, followed by two 1s, because the next two notes are the same as the previous note.

Example 4.5 shows an encoding for a four-note AD pattern descriptor pattern together with its musical notation. The pitches are the same as in Example 4.3 but the code shows that the second note is twice as long as the first, the third note is the same length as the second, and the fourth note is half the length of the third.

F+1:1/U0A:2/D2G+1:1/U0A:0.5

Example 4.5: Example of the encoding of a four-note AD pattern descriptor pattern type

4.4.5 Visual Basic Code

Once the database had been designed, it was necessary to transfer the MusicXML documents into the tables, which required a conversion program to be written. A number of different programming languages could have been employed to achieve this because most have the ability to read XML documents as well as the ability to connect to databases. As with the DBMS, there were no business-critical issues or issues surrounding distribution of the program to consider. Therefore, Microsoft's *Visual Basic 2010* was chosen because it was capable of accomplishing the task, it is a relatively straight forward language with short development times, and because I have used the language on other projects.

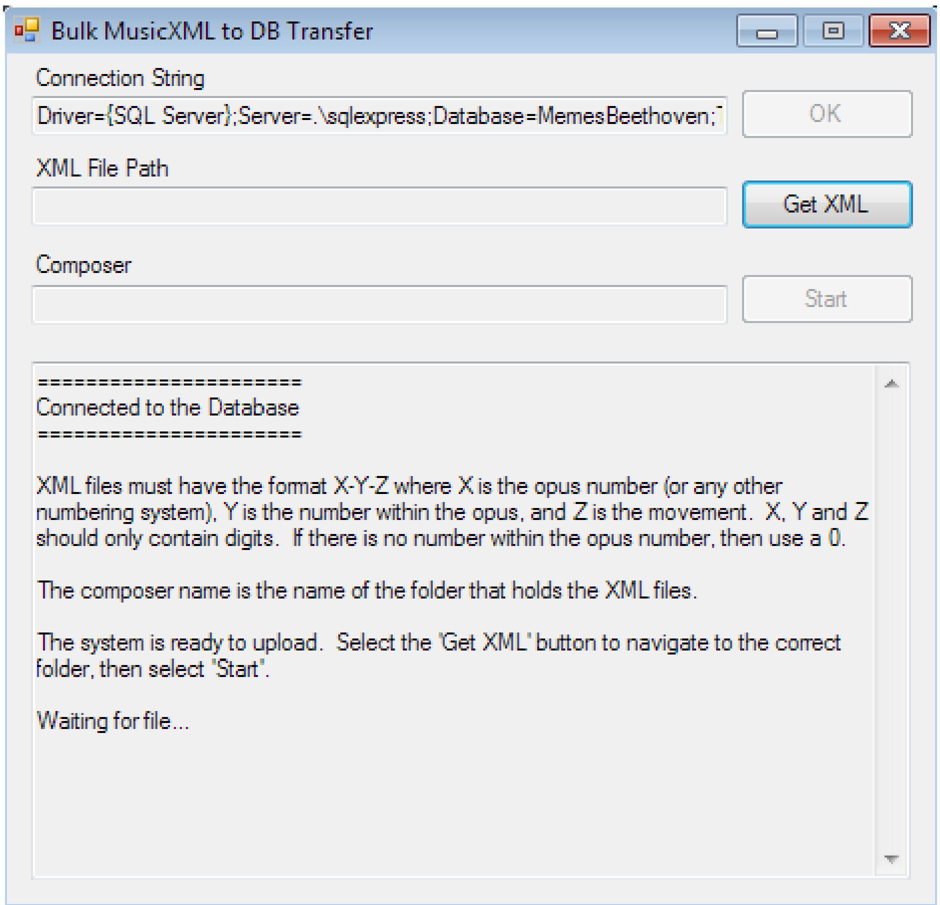

Figure 4.10: Visual Basic conversion program form

The *Visual Basic* project consists of one form (i.e., a single screen that permits a user to control the program), which allows the user to connect to the database and choose the MusicXML documents to be transferred, with the code needed to connect and transfer the data to the database being written as part of this form. The form consists of an editable field in which the database connection string is displayed, a non-editable field containing the chosen file path, a non-editable field containing the name of the composer whose works are being transferred, and a non-editable field showing the progress of the transfer and any system messages that occur. There are also three command buttons (buttons that invoke a piece of code): one that checks the connection string and connects the program to the database, one that allows the user via a dialog box to choose which folder to take the MusicXML documents from, and one that starts the process of transferring the data to the database (see Figure 4.10).

When the form is first opened, it tries to connect to the database. If it is unable to connect, it opens the form with a message in the box at the bottom and allows the user to amend the connection string. Once the connection string has been amended, the user then clicks on 'OK'. The program then tries again to connect to the database. If it is successful, instructions are displayed in the box at the bottom and the 'GetXML' button is enabled.

When the user clicks on the 'GetXML' button, a 'Browse for Folder' dialogue box is opened, which allows the user to navigate to a folder where the MusicXML documents are stored. Once this has been completed, the user can then click on 'Start'.

Clicking on the 'Start' button triggers a number of different processes. Firstly, the system checks to see if a record of the composer exists in the database. If the composer does not already exist in the database, the system creates a record for that composer. The system then loops through all the XML documents (each document contains one movement) within the chosen folder. For each XML document, it checks that the document has not already been uploaded; if it has, it ignores the document. If the document has not been uploaded already, then it looks to see if any other movements of the same piece have been uploaded. If any other movements have been uploaded then it just uploads the document. If no other movements have already been uploaded, then it creates a record for the piece before uploading the document (see Figure 4.11).

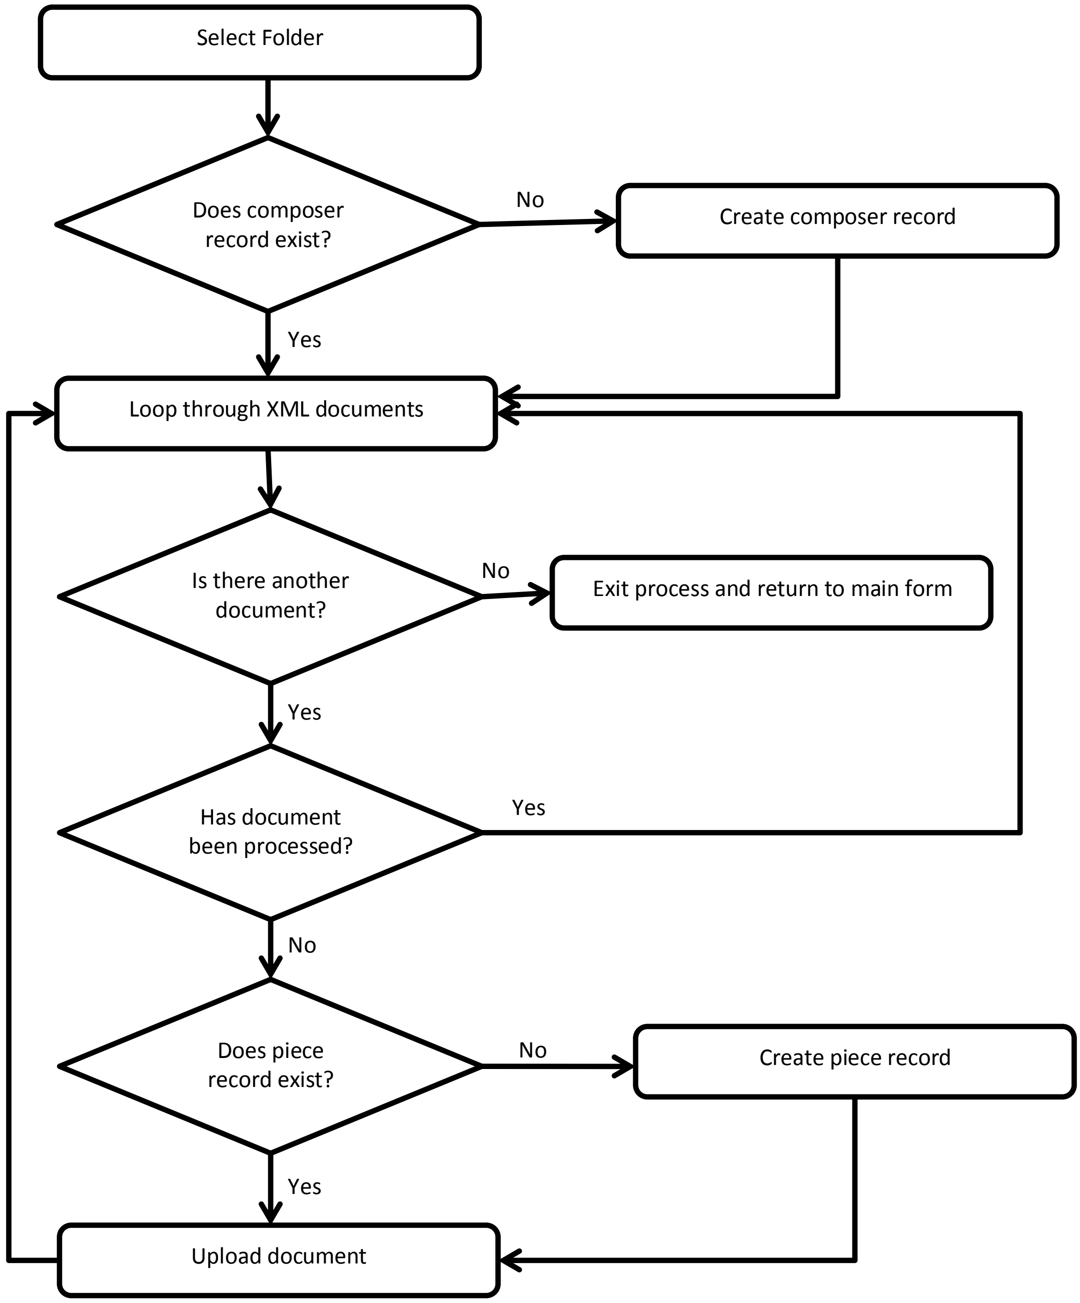

Figure 4.11: Flow chart for the uploading of XML documents to the database

The actual uploading of the document uses amended code from Deitel and Deitel (2009, p. 996). This code looks at the elements and attributes for each note and, for those elements and attributes that need uploading, it puts the values into variables.[4] When all the details for a note are obtained, it then adds the note to the database. When the code reaches the end of a document it moves onto the next

[4] A variable in coding terms is a temporary storage mechanism for data.

document, if there is one, or displays a completion message if all the documents in that folder have been processed.

4.5 Data Enrichment

The previous two sections have explained how the database is structured and how the MusicXML documents have been transferred into the tables. The following stage is to return to the KDD process with the next step of data enrichment.

Data enrichment concerns adding information that will help with the understanding of the data. There were two areas where information was added to the database: firstly, a table was used to help identify when notes were enharmonically equivalent and to help calculate intervallic distances between notes (tblNoteLookup); and, secondly, a table was used to group the year of each composition into periods (tblPeriod).

The table tblNoteLookup contains all possible pitches in all possible octaves for the instruments of a string quartet. All notes in a given octave that are enharmonically equivalent were assigned the same NoteLookupID. Each NoteLookupID is in sequence so that when the sounding pitch goes up by one semitone, the NoteLookupID also goes up one semitone. This table is then used to calculate the intervallic distance between two consecutive notes.

The table tblPeriod contains the PeriodID. Each PeriodID groups together a set of compositions written within a predetermined time-frame. This then allows comparisons to take place between pieces written around the same time with other pieces written in either an earlier or later time-frame. Creating periods for groups of compositions, rather than treating each composition as a separate entity, follows the concept of grouping similar items together as part of the KDD process. Each PeriodID covers a number of years and is allocated in ascending order, i.e., the earlier periods have a lower PeriodID than the later periods.

The periods were created after taking into account the repertoire that was imported into the database. However, before this could be done the year of each composition needed to be obtained and entered into the database.

There were a number of different factors to take into account when determining which year to count as the date of the composition. Each composition could have a number of different dates associated with it, such as the year the composition was started or finished, the year it was first published, or the year it was first performed. Other issues could have an impact: for example, there may have been sketches for the work from much earlier, or revisions could have been made much later.

For the sake of consistency, the year used was the year of composition as stated by *Grove Music Online* (2011). Where the period of composition spanned more than one year, the earliest year was used. This was decided on the grounds that most composers probably had the main ideas for the composition at the start of composing the work. However, it should be recognised that for memes to propagate other people need to hear or see the work, and this would not necessarily have taken place in the same year that the composer wrote the composition.

The periods fall roughly into five standard musicological periods: early Classical, late Classical, early Romantic, late Romantic, and early Twentieth Century. To begin, the first period covers the time from the first string quartet selected, Haydn's op. 1 no. 1, started in 1762, to Mozart's K. 173, started in 1773. The second period then starts with Haydn's op. 33 quartets, started in 1781. Because there was a break of eight years between Haydn's op. 20 quartets, started in 1772, and his op. 33 quartets, started in 1781, as well as a break between Mozart's K. 173, started in 1773, and his K. 387 quartet, started in 1782, this was where the break between the periods was placed. This means that the first period covers the early quartets of Haydn and Mozart, and the second period covers the later quartets of Haydn and Mozart as well as the first period quartets of Beethoven[5] (his op. 18, started in 1798).

The third period begins with Beethoven's op. 59 quartets, started in 1806, which are the first quartets to be written in Beethoven's middle period. This point also marks the end of all the Haydn and Mozart quartets used within the present research with the last Mozart quartet being his K. 590, started in 1790, and the last Haydn quartet being his op. 77 quartets[6], started in 1799, making a natural break between the second and third periods. The third period ends with Mendelssohn's op. 80, started in 1847. Therefore, this period covers the middle and late period quartets of Beethoven, together with quartets by Schubert, Mendelssohn and Schumann.

Brahms' op. 51 quartets, started in 1865, begin the fourth period, which provides a natural break from Mendelssohn's op. 80 quartet. The end of the fourth period is more problematic because there is not such a natural gap as with the other periods. This is because there are overlapping composers who are not naturally grouped together. For example, Debussy's op. 10 was started in 1893, the same year that Dvorak started his op. 96 quartet. Also, both Debussy's and Ravel's (1902) quartets were started before Sibelius' op. 56 quartet (1909). This leads to the question of whether Debussy and Ravel should be

[5] Although the standard three periods of Beethoven's output do not all conform to the periods selected for this study.
[6] Haydn's string quartet op. 103 was not included because it is incomplete.

included in the fourth period with Borodin, Brahms, Dvorak, Grieg, Smetana and Tchaikovsky, or whether they should be included with Bartok, Janacek, Prokofiev, Shostakovich and Sibelius in the fifth. In the end, the break was placed between Ravel's op. 35 quartet, started in 1902, and Bartok's op. 1 quartet, started in 1908, to keep the years in sequence and to separate Debussy and Ravel from Bartok (see Table 4.18).

Period	Years Covered	Composers Covered	Total Nº of Compositions	Total Nº of Movements
1	1762 – 1773	Haydn Mozart	25	97
2	1781 – 1799	Haydn Mozart Beethoven	44	177
3	1806 – 1847	Beethoven Schubert Mendelssohn Schumann	21	87
4	1865 – 1902	Smetana Borodin Brahms Tchaikovsky Dvorak Grieg Debussy Ravel	13	52
5	1908 – 1970	Janacek Sibelius Bartok Prokofiev Shostakovich	8	29

Table 4.18: Breakdown of the periods together with the associated composers within the dataset

4.6 Data Transformation

Data transformation concerns encoding the data in such a way that the amount of data stored is reduced or altered for ease of mining. For example, people who live in the same town could be grouped together, allowing analysis to take place by town rather than looking at each person individually. In the present research, a number of different areas of the data needed transforming from matching the compositions to the periods, to generating the patterns and their properties. The data transformation was implemented using 'stored procedures' (a series of sequential Transact SQL commands that are grouped together, which can then be run using one Transact SQL command).

4.6.1 Periods

The first stored procedure, Pass00:1Periods, was used to match the compositions with the periods. This was done by comparing the PieceYear column in tblPiece with the PeriodStart and PeriodEnd columns in tblPeriod and placing the appropriate value for the period in the PeriodID column of tblPiece.

4.6.2 Flagging Instrumental Parts

The next stage was to flag the start of pieces, movements and parts within tblNote so that the system knows when to stop generating patterns, and when to begin generating patterns for the next piece, movement or part. This was achieved by flagging the first note of each instrumental part by placing the number 1 in the NewPart column of tblNote.

A stored procedure called Pass01:1Parts looks for any notes that are in the first bar of the movement with a NoteOrder of 1 (i.e., the first note of bar 1) and places a 1 in the NewPart column. The stored procedure then places a 0 against all other notes to remove the NULL value from that column.[7]

4.6.3 Double Stopping

When there is double stopping on an instrument, the pattern generation system needs to know which notes from the chord to use. It is possible to create patterns that will use all of the notes in the chord by going through the chord(s) voice by voice. However, this approach would be complex in that the system would need to go through the chord(s) multiple times, taking each line created by the chord note individually, together with the appropriate number of notes before and after the chord(s). Because of time constraints, the simplest option was taken in that only one note from each chord was used in the pattern generation. Again for simplicity, it was always the top note from the violin and viola parts and the bottom note from the cello part.

A number of stored procedures were used to achieve this. Stored procedures Pass02:1Chords, Pass02:2Chords, Pass02:3Chords, and Pass02:4Chords for the instrumental parts violin 1, violin 2, viola and cello respectively looked for all the chords and flagged the notes to use in those chords

[7] Null values in databases are notoriously difficult to deal with. This is because a column entry can have three different states; it can have data, be Empty, or be Null. When the entry is Empty, the database knows what type of data the entry can hold (i.e., text, number, date, etc.). However, when the entry is Null, the database does *not* know what type of data the entry can hold. As a consequence, when querying the column it is necessary to be wary of Null values because they will not always be included in the results. Therefore, it is often easier to give the entries a value or to give them Empty status to minimise inaccuracies when querying the data.

using a 2 in the Chord column of tblNote. The next stored procedure, Pass02:5Chords, flags all notes that are not part of a chord using a 1 in the Chord column of tblNote. Finally, stored procedure Pass02:6Chords looks at all the remaining notes (i.e., those that are part of a chord but are not being used in the pattern generation) and flags these notes using a 3 in the Chord column of tblNote.

4.6.4 Tied Notes

Because only pitches and durations are being considered in the present research, any tied notes can be combined into one single note with their durations being summed.

The stored procedure Pass03:1Ties goes through tblNote looking for any notes that are tied together using the Tie column in tblNote, and also for any consecutive rests. When it finds such notes, or rests, it adds the durations of all the tied notes, or those of the rests, together and replaces the duration of the last note of the tie, or the last rest of the group of rests, with this calculated value. It then removes all the other notes connected by that tie or removes all other the rests within that group of rests. For example, two crotchets tied together become a single minim, and a crochet rest followed by a quaver rest becomes a single dotted crotchet rest.

As explained in Section 4.3 above, it was not always possible to distinguish between ties and slurs within *Photoscore*. Because there may have been some discrepancies in the data, the stored procedure will only treat two consecutive notes as a tie if they are the same pitch (including their enharmonic equivalent), even if the data in the Tie column of tblNote indicates the notes are tied.

4.6.5 Absolute, Relative and Duration Calculation

In order to generate the patterns, the system needs to know the position of each note relative to its neighbours; i.e., when the system encounters an A, is that A higher or lower than the previous and subsequent notes? This is achieved by creating an absolute pitch value, a relative intervallic value, and a relative durational value for each note (corresponding to AbsValue, RelValue, and TimeValue respectively in tblNote) in relation to the previous note.

The stored procedure Pass04:1NoteValues goes through each note in tblNote and uses an encoding system (as described in Section 4.4.4 above) to show the position in pitch and duration of a note relative to its predecessor note. Because the system looks at both absolute and relative intervallic values for pitches, as well as including duration information, the system makes three different calculations for each note.

Firstly, for each record of a note in tblNote (except for the first note of a piece, movement or part) the stored procedure works out the relative position of the absolute pitch value of the note by calculating whether the note is higher or lower than the preceding note (disregarding any rests that separate the notes). If the note is the same pitch, it puts an 'S' into the AbsValue column of tblNote. Likewise, if the pitch is lower than the previous note it uses a 'D' and if the pitch is higher than the previous note it uses a 'U'. It then calculates how many octaves are between the two notes and places that value next to the direction indicator in the AbsValue column of tblNote. Finally, it places the actual pitch of the note next to the octave counter in the AbsValue column of tblNote. For example, if the first pitch is an F in octave 3 and the second pitch is a G sharp in octave 4 it will place 'U1G+1' in the AbsValue of tblNote of the G sharp record, indicating that the second pitch is going up by at least an octave to a G sharp from the first pitch. When the record in tblNote represents a rest, '00R' is placed in the AbsValue column of tblNote because no direction, pitch, or octave placement is required with rests.

Secondly, the stored procedure takes each record of a note in tblNote (except for the first note of a piece, movement or part) and works out the relative intervallic distance of the note by calculating whether it is higher or lower than the preceding note (disregarding any rests that separate the notes), and then works out how many semitones there are between the two notes (using tblNoteLookup). Like the absolute pitch value calculation, if the note is the same pitch it puts an 'S' into the RelValue column of tblNote, if the pitch is lower than the previous note it uses a 'D', and if the pitch is higher than the previous note it uses a 'U'. It then calculates how many semitones there are between the two notes (using tblNoteLookup) and places that value next to the direction indicator in the RelValue column of tblNote. For example, if the first pitch is an F in octave 3 and the second pitch is a G sharp in octave 4 it will place 'U15' in the RelValue of tblNote indicating that the second pitch is 15 semitones higher than the first pitch. When the record in tblNote represents a rest, 'R0' is placed in the RelValue column of tblNote.

Finally, the stored procedure takes the duration of each note or rest record and calculates its relationship to the preceding note or rest record's duration. If the previous record has the same duration it places a '1' in the TimeValue column of tblNote indicating that both records have the same duration. If the previous record is not the same value then it divides the duration of the second record by that of the first record and places that value in the TimeValue column of tblNote. This means that if the second note is longer than the first, the TimeValue will be greater than 1. For example, if the duration of the first note is 4 and the duration of the second is 8, the TimeValue will be 8 divided by 4 giving a TimeValue of 2. This effectively means that the second note is twice as long as the first. Conversely, if the second note is shorter than the first, the TimeValue will be less than 1. For example, if the duration of the first note is 8 and the duration of the second is 4 the TimeValue will be 4 divided by 8 making a TimeValue of 0.5. This effectively means that the second note is half the length of the first. Unlike the absolute and relative calculations, rests are not treated any differently from notes, for both have duration.

4.6.6 Pattern Generation

Once the previous phases have been completed, the data is finally ready for the system to generate the patterns. Patterns are generated for all possible combinations of three- to eleven-note consecutive single-line patterns within each piece, as is consistent with the definition of a pattern in music as described in Section 2.5.3 above, and these are placed in tblPatterns. There are three stored procedures that perform this task: Pass05:1Patterns, Pass05:2Patterns, and Pass05:3Patterns.

The first stored procedure, Pass05:1Patterns, begins by disabling the indexes on tblPatterns to speed up the process. Indexes are rebuilt each time a record is added and because this procedure adds the generated patterns (a generated pattern is equivalent to one record) one at a time, it drastically increases the time taken for this process. Therefore, the indexes were disabled to stop this process and were enabled again at the end of the process in Pass05:3Patterns. The re-enabling of indexes automatically rebuilds them, meaning that nothing has been lost in the process of disabling and re-enabling the indexes.

Once the indexes have been disabled, the patterns are then generated. The basic process for pattern generation is as follows:

> **Step 1** – Extract the first eleven notes of an instrumental part and place the AbsValue, RelValue and TimeValue for each of these notes into variables.
>
> **Step 2** – Change the first note variable for AbsValue to just the pitch name, the RelValue to 0, and the TimeValue to 1 (this is because these values are for the first note in the pattern and therefore do not need to be related to a previous note).
>
> **Step 3** – Take the first three sets of values in the variables and create a three-note pattern putting the resultant pattern into tblPatterns.
>
> **Step 4** – Repeat step 3 for all the four- to eleven-note patterns.
>
> **Step 5** – Move all the variables down one (i.e., move the second note's AbsValue into the first note's AbsValue variable, the third note's AbsValue into the second note's AbsValue, and so on).

Step 6 – Add the next note from tblNote into the eleventh note's variables

Step 7 – Repeat the process from **Step 2**.

Example 4.6 gives an example of patterns generated by this algorithm using a short six-note pattern.

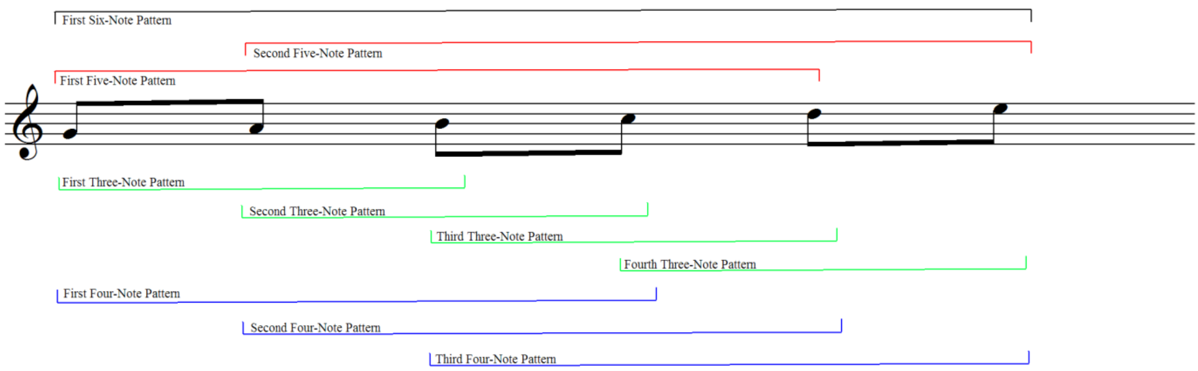

Example 4.6: An example of pattern generation

The only exception to this algorithm is when the system reaches the end of a piece, movement or part. When this happens, it keeps moving the variables down without adding any more notes from tblNote and creates the appropriate length patterns whilst reducing the number of notes each time until it has processed the last three-note pattern (i.e., it repeats **Step 2** to **Step 7** ignoring **Step 6**). Then it begins from **Step 1** again with the next appropriate piece, movement or part.

The next stored procedure Pass05:2Patterns makes two modifications to the records in tblPatterns. First, it removes any patterns from tblPatterns that have a rest as either the first or last element of the pattern. With this system there is no recognisable difference between a four-note pattern of rest, A, G sharp, A and a three-note pattern of A, G sharp, A. Secondly, the stored procedure removes any null values in tblPatterns.

Finally, stored procedure Pass05:3Patterns re-enables the indexes on tblPatterns to speed up the next stage of the system.

4.6.7 Pattern Properties

The final stage of data transformation is to identify some basic properties for patterns that will allow a comparison of similarity between different patterns to take place according to the criteria laid out in Section 3.3 above. The patterns in tblPatterns are grouped together according to whether they are the same in terms of their absolute pitches, their relative intervallic values, or their relative durations. For example, when looking at the absolute values of pitches, all patterns comprising of just three Cs in the same octave will be grouped together. When looking at the relative intervallic values, all patterns that have three notes at the same pitch in the same octave will be grouped together. For each group of patterns, a record is created in tblPatternProperty together with the properties for that particular pattern. Reference is then made back to tblPatterns by placing the PatternPropertyID from tblPatternProperty into the appropriate column of tblPatterns (AbsID, RelID, or TimeID).

The pattern properties are worked out as follows:

Number of Notes – how many notes are in the pattern

Pitch Shape – the overall direction between the pitches (i.e., a combination of ups, downs, or same notes).

Duration Shape – the overall movement between the durations (i.e., a combination of longer, shorter, or same length durations).

Pitch Centre – adds all the relative intervallic distance values together, then divides by the length of the pattern minus 1.

Duration Centre – Adds all the relative durational values together, then divides by the length of the pattern.

Pitch High/Low – calculates the intervallic distance between the highest and lowest notes

Duration High/Low – calculates the distance between the longest and shortest notes

The pattern properties are calculated using a series of stored procedures beginning with Pass06:0Index. Like the generation of patterns, the first task is for the system to disable the indexes on tblPatternProperty to speed up the process. This is done by Pass06:0Index.

Next, the stored procedures Pass06:1:1Time, Pass06:2:1Rel, and Pass06:3:1Abs group the same patterns from tblPatterns together, place a record for each distinct pattern in tblPatternProperty, and fill in the PropertyType, NumNotes, and Pattern columns of tblPatternProperty. Stored procedures Pass06:1:2Time, Pass06:2:2Rel, and Pass06:3:2Abs update the TimeID, RelID and AbsID respectively in tblPatterns with the appropriate PatternPropertyID from tblPatternProperty. Stored procedures Pass06:1:3Time, Pass06:2:3Rel, and Pass06:3:3Abs calculate the appropriate pitch or duration shape of the patterns and their pitch or time centre value, and place the results into the Shape and PitchCentre (stores both the duration and pitch centre values) columns of tblPatternProperty respectively. Then, stored procedures Pass06:1:4Time, Pass06:2:4Rel, and Pass06:3:4Abs calculate the difference between the highest and lowest value of the pattern and place this value in the HighLow column of tblPatternProperty. Finally, the indexes on tblPatternProperty are re-enabled by stored procedure Pass06:4Index.

4.7 Data Mining

Once the previous four stages of KDD have taken place, the next stage of data mining, which involves processing and analysing the data, can take place. Elmasri and Navathe list five different categories of information that can be gained during data mining: *Association rules* – if one item is present then what is the probability of a different item being present; *classification hierarchies* – items are grouped together with similar items under a hierarchical structure; *sequential patterns* – if one event occurs what is the likelihood of another event occurring later; *patterns within time series* – looking at items of data across a time series; and *clustering* – do similar items cluster at certain points within the dataset (2007, pp. 948-949). The present research looks at patterns within music across a number of different time-frames, therefore the data mining will aim to produce data that relates to Elmasri and Navathe's 'patterns within time series' categorisation.

There are a number of different algorithms that can be used for the data mining process, such as Sampling Algorithms, Apriori Algorithms, Frequent Pattern Tree Algorithms, Partition Algorithms, etc. (Elmasri & Navathe, 2007, pp. 947-957). All of these algorithms are based on the idea that certain items within the dataset can be removed due to their infrequent occurrences, which helps speed up the processing. Because the present research aims to find infrequent patterns as well as frequent patterns within the dataset, none of the algorithms introduced by Elmasri and Navathe is appropriate. However, the Association Rules algorithm (i.e., what is the likelihood of one item being connected to another item) can be adapted (removing the *Support* and *Confidence* computations, which are there to reduce the quantity of data being mined), because it basically ranks recurring item sets (i.e., in the present research, patterns of notes) according to the frequency of their occurrence.

For the present research, the data mining involves counting the patterns, ranking them according to their frequencies, and creating a table that shows each pattern's progression across the ranking positions over time (tblRanking). The stored procedures Pass07:1Ranking and Pass07:2Ranking calculate the ranking position for each pattern within each period with Pass07:1Ranking creating the ranking position across the different pattern lengths and Pass07:2Ranking creating the ranking position within the same pattern length group.

The ‘frequency’ of a pattern is equal to the number of times a pattern appears altogether within a period, multiplied by the number of movements it appears in within that period. For example, if a pattern appears 100 times within a period across 4 movements it has a Frequency of 400. However, if a pattern appears 50 times within a period across 10 movements it has a Frequency of 500. The result of this calculation is that those patterns appearing many times in many different movements will have a high frequency, and will consequently appear towards the top of the ranking positions. However, a pattern type that appears few times in one movement, or a pattern type that appears just once in many different movements will have a low frequency, and will consequently appear towards the bottom of the ranking positions. This is to stop a pattern that appears many times but only in a small number of movements from gaining a high ranking position, as the potential propagation of the pattern is less owing to it appearing in a small number of works (further discussion on this process takes place in Section 5.5).

The first stored procedure (Pass07:1Ranking) ranks the patterns in descending order in each period according to their Frequency and places the results in tblRanking using the RankAll column for the ranking position. Stored procedure Pass07:2Ranking ranks the patterns within each group of pattern lengths in each period and updates the RankNum column in tblRanking with these results. Finally, the stored procedure Pass07:3Ranking puts the Shape value from tblPatternProperty into the PitchShape and DurationShape columns of tblRanking, calculates the pattern type encoding, and places this in the Pattern column of tblRanking.

4.8 Data Reporting

The final phase in KDD is to produce reports on the dataset. The purpose of reporting in the context of the present research was to be able to investigate the ranking positions of the patterns in order to determine whether there is any evidence for the evolutionary processes of selection, replication and variation, as well as the replicator properties of longevity, fecundity and copying-fidelity, as hypothesised in the working definition of a meme in music.

A number of stored procedures were written in order to extract the required information from the dataset. Each of these stored procedures produces information for one or more of the tables in the next three chapters. The data was either copied directly from the running of the stored procedure, or it was placed into Microsoft’s *Excel 2010* for further manipulation (such as calculating percentage figures). This information is discussed in Part III

4.9 Testing Phase

Three different sets of testing were undertaken to ensure the coding was accurate and the results produced would be what was expected. The first tests involved testing the code using two methods: putting a very small sample of data through at each stage of the implementation of the algorithm and manually verifying the results, and putting a larger subset of live data through each stage of the implementation and selecting some resultant records to manually check the accuracy of the code. The second stage consisted of sending two sets of data (a pseudo-random sample and the same pseudo-random sample seeded with three pre-determined eleven-note patterns) through the system to check the algorithms produced quantifiable results. Finally, the live data was put through each stage of the algorithm and randomly selected records were again checked to verify the accuracy of the coding.

4.9.1 Code Testing

The Visual Basic code was tested firstly by using a manually created XML document containing a small number of pitches and rhythms. This XML document was uploaded to a test database and each note was checked to verify that it had been encoded correctly within the database. A small selection of XML documents from the live dataset were then uploaded to the test database. The start and end note of each instrumental part and movement were then checked to see that they had been correctly encoded within the test dataset. Secondly, a random set of notes were selected from the test dataset and checked against an original printed score to ensure the accuracy of the encoding.

Each stored procedure within the database was tested using both a small set of pre-determined data and a set of selected records from a sub-set of the live data. The pre-determined data was designed at each stage to ensure that the data was appropriate for the stored procedure being tested. For example, when testing the stored procedure Pass01:1Parts, records were created in the pre-determined data for each break between instrumental parts, movements, and pieces. These breaks were then manually checked to verify the stored procedure had correctly identified the start and end of each instrumental part, movement, or piece. A subset of the live dataset was then passed through each of the stored procedures. After each stored procedure was run, a set of records were selected according to the expected outcome of the procedure to verify if the procedure had run correctly. For example, after the stored procedure Pass01:1Parts was run, a number of the records at the start and end of each instrumental part, movement and piece were checked to see if they had been flagged correctly.

4.9.2 Algorithm Testing

The first test involved creating a pseudo-random sample of pitches and passing them through the algorithm. With pseudo-random pitches it would be expected that all patterns would have the same ranking position, providing that there are no duplicate patterns within the dataset. A stored procedure was created within a test database to provide 12,000 pseudo-random pitches (approximately 1% of the size of the live dataset) from 47 different pitches (the average span of pitches for the instrumental parts within the live dataset) with four different pitch durations (although there are only four pitch durations, due to the pseudo-random selection of the durations, some of the resultant rhythms can become complex, e.g., a crotchet followed by a quaver followed by a semi-quaver followed by a crotchet). The pseudo-random pitches were split into five periods (the same number as the live dataset), with two pieces both containing four movements per period (to ensure a spread of patterns across both pieces and movements), resulting in each of the four instrumental parts containing 75 notes. Table 4.19 shows the resultant number of ranking positions for each PTD/PTL/Period[8] combination.

Table 4.19 clearly shows that there is only one ranking position for the vast majority of PTD/PTL/Period combinations. For the shorter-length PTDs (such as the three- and four-note PTLs) there is sometimes more than one ranking position, especially for the R-PTs. This difference in the number of ranking positions between the shorter- and longer-length PTDs is to be expected as there are more possible combinations of notes for the longer- than the shorter-length PTDs. For example, a three-note A-PT has 48 x 48 x 48 (i.e. 48^3)[9] possible permutations whereas an eleven-note A-PT has 48^{11} permutations. Additionally, the R-PTs have less possible permutations per PTL than their A-PT counterparts. For example, a three-note A-PT has 48^3 possible permutations whereas a three-note R-PT only has 48^2 possible permutations (there are only two possible choices as it does not matter which note the permutation begins with under the R-PTs). Altogether there are 2,400 notes in each period creating 2,384[10] three-note patterns, and there are 2,304 (i.e., 48^2) possible three-note R-PTs, meaning that there will always be some duplicate patterns generated for three-note R-PTs using a pseudo-random dataset of 12,000 notes. For patterns of four or more notes, it is possible that there will be no duplicate patterns.

[8] A description of PTDs and PTLs is provided in Section 5.1.
[9] 47 pitches plus a rest.
[10] When taking into account the breaks between pieces/movements/instrumental parts.

		Period				
PTD	PTL	1	2	3	4	5
A	3	3	4	3	3	3
	4	2	1	3	2	1
	5	1	1	1	1	1
	6	1	1	1	1	1
	7	1	1	1	1	1
	8	1	1	1	1	1
	9	1	1	1	1	1
	10	1	1	1	1	1
	11	1	1	1	1	1
AD	3	2	2	1	1	2
	4	1	1	1	1	1
	5	1	1	1	1	1
	6	1	1	1	1	1
	7	1	1	1	1	1
	8	1	1	1	1	1
	9	1	1	1	1	1
	10	1	1	1	1	1
	11	1	1	1	1	1
R	3	7	9	7	9	8
	4	2	3	3	2	2
	5	1	1	1	1	1
	6	1	1	1	1	1
	7	1	1	1	1	1
	8	1	1	1	1	1
	9	1	1	1	1	1
	10	1	1	1	1	1
	11	1	1	1	1	1
RD	3	3	2	2	2	3
	4	1	2	1	1	1
	5	1	1	1	1	1
	6	1	1	1	1	1
	7	1	1	1	1	1
	8	1	1	1	1	1
	9	1	1	1	1	1
	10	1	1	1	1	1
	11	1	1	1	1	1

Table 4.19: Number of ranking positions for each PTD/PTL/Period combination using the pseudo-random dataset

The second test involved seeding the pseudo-random dataset with three pre-determined eleven-note patterns. The ranking positions of these seeded patterns should appear higher in the ranking positions than those patterns generated from the pseudo-random data. Example 4.7 shows the three eleven-note patterns used to seed the pseudo-random dataset. Each pattern in Example 4.7 was created so that there are no duplicates of any three- to eleven-note length patterns across all three examples, or within each example.

i) First eleven-note seeded pattern

ii) Second eleven-note seeded pattern

ii) Third eleven-note seeded pattern

Example 4.7: The three eleven-note patterns used to seed the pseudo-random dataset

Each pattern in Example 4.7 was then placed at random within the pseudo-random dataset with a different number of appearances within each period as stated in Table 4.20.

Example	Period				
	1	2	3	4	5
i	16	12	8	4	0
ii	0	4	8	12	16
iii	8	8	8	8	8

Table 4.20: The number of times each pattern in Example 4.7 appears within each period

The patterns used in Example 4.7 have an effect on the number of ranking positions for each PTD/PTL/Period combinations compared to the purely pseudo-random dataset. Table 4.21 shows the number of ranking positions for each PTD/PTL/Period combination using the seeded pseudo-random dataset.

		Period				
PTD	PTL	1	2	3	4	5
A	3	5	7	7	7	5
	4	5	7	7	7	5
	5	5	7	7	7	5
	6	5	7	7	7	5
	7	5	7	7	7	5
	8	5	7	7	7	5
	9	5	7	7	7	5
	10	5	7	7	7	5
	11	5	7	7	7	5
AD	3	4	4	2	5	5
	4	4	4	2	5	5
	5	4	4	2	5	5
	6	4	4	2	5	5
	7	4	4	2	5	5
	8	4	4	2	5	5
	9	4	4	2	5	5
	10	4	4	2	5	5
	11	4	4	2	5	5
R	3	15	16	15	21	19
	4	15	16	15	21	19
	5	15	16	15	21	19
	6	15	16	15	21	19
	7	15	16	15	21	19
	8	15	16	15	21	19
	9	15	16	15	21	19
	10	15	16	15	21	19
	11	15	16	15	21	19
RD	3	6	7	5	9	6
	4	6	7	5	9	6
	5	6	7	5	9	6
	6	6	7	5	9	6
	7	6	7	5	9	6
	8	6	7	5	9	6
	9	6	7	5	9	6
	10	6	7	5	9	6
	11	6	7	5	9	6

Table 4.21: Number of ranking positions for each PTD/PTL/Period combination using the seeded pseudo-random dataset

Table 4.21 shows that the seeded pseudo-random dataset produces a greater number of ranking positions for all the PTD/PTL/Period combinations than the non-seeded pseudo-random dataset in Table 4.19. This is what would be expected as the seeded patterns are inserted multiple times into the pseudo-random dataset giving them a greater probability of a higher frequency figure. To show whether the seeded patterns are rising or falling through the ranking positions across the periods, Table 4.22 shows the ranking position of the first pattern of the three- to eleven-note patterns generated from the seeded patterns (i.e., only those patterns generated by the system that start with the first note of the three seeded patterns).

		Seed 1					Seed 2					Seed 3				
	Period	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
PTD	PTL															
A	3	1	1	3	4	-	-	4	2	1	1	2	2	3	2	2
	4	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	5	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	6	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	7	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	8	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	9	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	10	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	11	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
AD	3	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	4	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	5	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	6	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	7	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	8	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	9	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	10	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	11	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
R	3	4	4	7	15	19	12	12	6	3	1	1	1	1	2	9
	4	1	1	2	4	-	-	3	2	1	1	3	2	2	2	2
	5	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	6	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	7	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	8	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	9	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	10	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	11	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
RD	3	2	2	3	6	-	6	5	3	2	1	1	1	1	3	3
	4	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	5	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	6	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	7	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	8	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	9	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	10	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2
	11	1	1	1	3	-	-	3	1	1	1	2	2	1	2	2

Table 4.22: The ranking position in each period of the first pattern generated from each of the seeded patterns within the pseudo-random dataset

We would expect that the ranking positions for each of the seeded patterns would behave as shown in Table 4.23. This is because the frequency with which each seeded pattern was placed within the pseudo-random data should mean that each seeded pattern will have a greater number of appearances within the appropriate period than any of the pseudo-random generated patterns. Additionally, each seeded pattern's frequency has been set so that they follow a certain path through the periods in terms of their ranking positions.

Seeded Pattern	Period				
	1	2	3	4	5
Example 4.7 i	1	1	1	3	
Example 4.7 ii		3	1	1	1
Example 4.7 iii	2	2	1	2	2

Table 4.23: Expected ranking position for each of the seeded patterns if the pseudo-random data has no impact

As can be seen from Table 4.22, the vast majority of the PTD/PTL combination seeded patterns follow the ranking positions set out in Table 4.22. The remaining PTD/PTL combination seeded patterns ranking positions can all be attributed to interference from the pseudo-random dataset. For example, the unexpected three-note A-PT ranking position in periods 3 and 4 is a consequence of other three-note patterns generated from the seeded patterns having duplicates within the pseudo-random dataset. This means that the three-note patterns generated from the seeded patterns that have duplicates within the pseudo-random dataset will have a greater frequency of occurrences than those that have no duplicates, giving the patterns with the duplicates a higher ranking position.

However, the amount of interference will be negligible in terms of the present research. The figures for the number of possible distinct patterns that can be created from a random generation system are, apart from some three- and four-note PTDs, less than the total number of patterns within the live dataset. Where interference exists it will mainly affect the bottom ranking positions due to the number of patterns within the dataset that are actually repeated, as shown by the position of the seeded patterns within the ranking positions in Table 4.22. Because the majority of the analysis is looking at the movement across the ranking positions of the patterns, the patterns affected by interference will still need to progress to or from the top ranking positions, and consequently will still need to show a large amount of duplication within some periods compared to other periods.

4.9.3 Live Data Testing

After all the data was uploaded from the XML documents into the live database, a sample of records was selected and checked against a printed score to verify that the uploading was accurate. Furthermore, after each stored procedure had been run, a set of relevant records were manually checked to ensure that each stage of the algorithm was being implemented correctly. For example, once stored procedure Pass02:1Chords had been run, a selection of records were chosen where the first violin part had either double stopping or more than one line. Each of these selected records was checked in tblNotes to ensure that the relevant record had been updated with the appropriate encoding in the Chord column.

4.10 Summary

Once the technologies had been alighted upon, the next stage was to demonstrate how the technologies could be used within the context of the present research. A description was given in this chapter of how the six phases of KDD, as proposed by Elmasri and Navathe (2007, pp. 946-947), were implemented within the present research. Firstly, data selection showed how the repertoire for the dataset was compiled using freely available existing MusicXML documents, converting freely available pdf documents, and scanning in repertoire to produce a set of compositions that covered a range of composers and periods. Secondly, data cleansing showed some of the difficulties in using certain technologies, such as inaccurate conversions, and in relying on freely available scores where errors such as incorrect notes were detected.

A diversion was then made from the concept of KDD to show how the database itself was designed. This included details of all the tables used and how they linked together to form a coherent dataset. Additionally, an explanation was given of how the MusicXML documents were transferred into the database using a conversion program created in *Visual Basic*.

Returning back to KDD, the next stage of data enrichment showed how enharmonically equivalent notes were dealt with in the system, and how the time-periods which grouped compositions from the same era were defined. Subsequently, it was shown how the stored procedures processed different transformations: how ties and chords were navigated by the system, how the absolute pitch, relative intervallic, and relative duration values were calculated, how the patterns were generated, and how the properties for the matching of similar patterns were generated. Following on from data transformation, data mining showed how the ranking positions were created using a frequency calculation based on the number of instances of a pattern multiplied by the number of movements in which the pattern occurred. Finally, data reporting showed how stored procedures were used to extract data that could be used to investigate the evolutionary processes, as well as the replicator properties.

The final section of this chapter looked at how the code and methodology was tested. The code was tested by manually checking samples of data after each stage of the implementation of the algorithms. The methodology was tested by using a pseudo-random dataset and a seeded pseudo-random dataset, both of which showed that the methodology performed as would be expected.

5 Chapter 5: Initial Observations

5.1 Introduction

The most striking aspect of the initial observations is the predominance at the top of the ranking position tables of three-note PTs that do not move from their starting pitch (i.e., same-note PTs). This was true for all five periods as well as for all four PTDs. For example, the highest ranking position A-PT in the final period was three consecutive Ds as in Example 5.1: Three consecutive Ds - The highest ranking position A-PT in Period 5 (n.b., the three consecutive Ds can be in any octave playable by the instruments of a string quartet).

Example 5.1: Three consecutive Ds - The highest ranking position A-PT in Period 5

However, is this predominance in the highest ranking position of three-note, same-note PTs a reflection of the most frequent PTs as a whole?

This chapter investigates the ranking positions in terms of some initial findings. Overall, it finds that the most prominent patterns across all the periods consist of three consecutive notes the same. When these same-note patterns are investigated further, it is discovered that the methodology allowed repeated patterns (such as a repeated bass note, as in the opening two bars of the 1st movement from Haydn's String Quartet op. 54 no. 1) to become more prominent. Additionally, three-note patterns are naturally more dominant than their longer-length counterparts owing to the number of different possible permutations of patterns each pattern length can generate. These issues are resolved by implementing a de-duplication procedure that only counts the first instance of a pattern in a repeated pattern of notes when calculating the frequency of a pattern, and by not comparing patterns of different lengths when calculating the ranking positions.

A number of terms and abbreviations are used to help with the understanding of this, and the following two chapters. Firstly, the present research looks for evidence for memes using absolute pitch values and relative intervallic values, both with or without relative durational values (as described in Section 4.4.4). These four views of the data are described as pattern type descriptors (PTDs). Within each pattern type descriptor there are a number of different pattern types (PTs), which describe the pitches and/or durations of one or more pattern type instances (PTIs). A pattern type instance refers to an actual appearance of a pattern type within the database. For convenience, the pattern types use the following abbreviations: A-PT for absolute pitch values without any duration, AD-PT for absolute pitch values with relative durations, R-PT for relative intervallic values without any duration, and RD-PT for relative intervallic values with relative durations. Finally, pattern type length is abbreviated to PTL. For reference, the abbreviations are listed in the header for the subsequent pages of both this chapter and the following two chapters.

5.2 Same-Note Pattern Types

Further investigation of the top of the ranking positions revealed that there was predominance amongst the 20 highest ranking positions of same note PTs. This occurred across all four PTDs and across all five periods, as shown in Table 5.1.

	Period									
	1		2		3		4		5	
PTD	Highest Ranked Same-Note PT	Nº in Top 20	Highest Ranked Same-Note PT	Nº in Top 20	Highest Ranked Same-Note PT	Nº in Top 20	Highest Ranked Same-Note PT	Nº in Top 20	Highest Ranked Same-Note PT	Nº in Top 20
A-PT	1	16	1	14	1	19	1	20	1	9
AD-PT	1	20	1	18	1	18	1	19	1	10
R-PT	1	7	1	7	1	8	1	7	1	6
RD-PT	1	9	1	8	1	8	1	7	1	5

Table 5.1: Predominance of same-note PTs in the 20 highest ranking positions

Table 5.1 shows both the ranking position of the highest ranked same-note PT and the number of same-note PTs there are in the 20 highest ranking positions for each period. It can be seen that the predominance of same-note PTs occurs across all five periods, although in Period 5 there is a distinct reduction in the number appearing in the 20 highest ranking positions, especially for R-PTs and RD-PTs.

A shift in the highest ranking position PTs would be expected as composers turned away from the tonal-centric music of the classical and romantic eras and began investigating alternative approaches to pitch organisation. However, same-note PTs within a composition may not always imply a tonal centre. Additionally, there are a number of accompaniment figures that use repeated notes that could have been rejected by twentieth-century composers and this would probably have an effect on the highest ranking position PTs. Nevertheless, it is relatively easy to find examples of a repeated-note accompaniment figure in twentieth-century string quartets (such as in the second violin part from bar 8 of the second movement of Bartok's String Quartet no. 2, and both the violin and viola parts from bar 139 of the fifth movement of Shostakovich's String Quartet no. 3).

A number of string quartets use a repeated note as an accompaniment figure indicated using the tremolo notation (for example, the opening of the first movement of Brahms' String Quartet op. 51 no. 1 in the viola and cello parts). As discussed in the previous chapter, owing to the complex nature of converting tremolos into MusicXML, these indicators were not included, nor were they expanded out into possible realisations, thereby reducing the number of same-note PTs generated. As such, same-note PTs are arguably under represented, making their dominance in the ranking positions unexpected.

It could be that the predominance of same-note PTs is a result of the methodology used to generate the patterns. The system took each instrumental part from a movement and created all the possible three- to eleven-note patterns. However, this created the situation of a group of notes all on the same pitch being counted multiple times. For example, a group of five consecutive Gs resulted in one five-note pattern, two four-note patterns, and three three-note patterns being generated (see Example 5.2). The advantage of this method is that it ensured that all possible patterns were generated and therefore all had the potential to exist as memes. Nevertheless, this method allowed for continuously repeated notes and patterns to attain more prominence than non-repeated groups of notes.

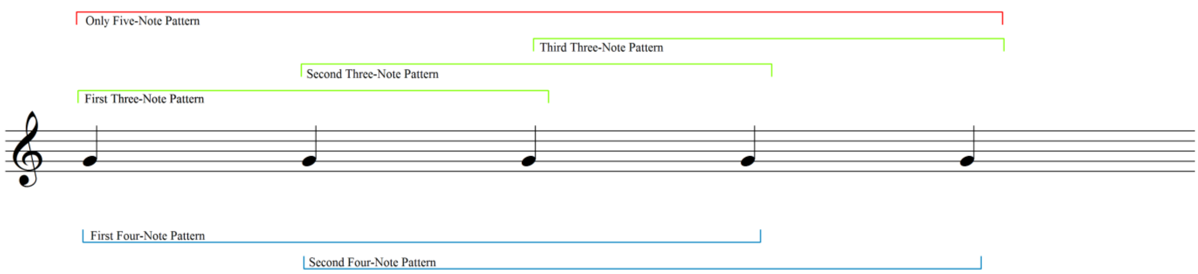

Example 5.2: Five consecutive notes pattern generation

5.3 The De-Duplication Process

In order to reduce the effect of the methodology on the multiplication of same-note PTs, a new set of stored procedures (Pass08:1DeDupA, Pass08:2DeDupR, Pass08:3DeDupAD, and Pass08:4DeDupRD) was created to count the repeated group of notes only once. For example, in Example 5.2 above, the five Gs under the new stored procedure produces only one five-note pattern, one four-note pattern, and one three-note pattern, as shown in Example 5.3.

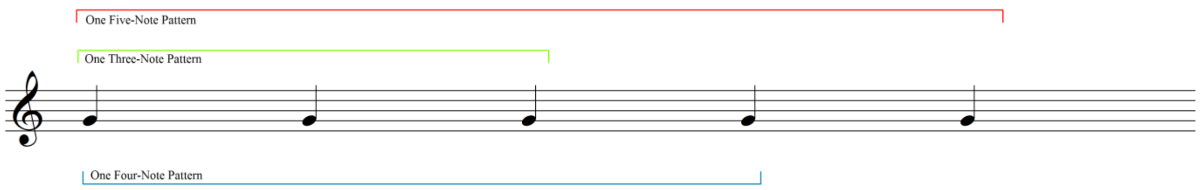

Example 5.3: Five consecutive notes pattern generation using de-duplication stored procedures

The new stored procedures also had an effect on all repeated figures, such as alternating pattern groups. For example, under the original pattern creation procedure a group of six notes alternating between a G and a D would generate one six-note pattern, two five-note patterns, three four-note patterns and four three-note patterns (see Example 5.4). But under the new stored procedure the system creates one six-note pattern, two five-note patterns, two four-note patterns, and two three-note patterns (see Example 5.5).

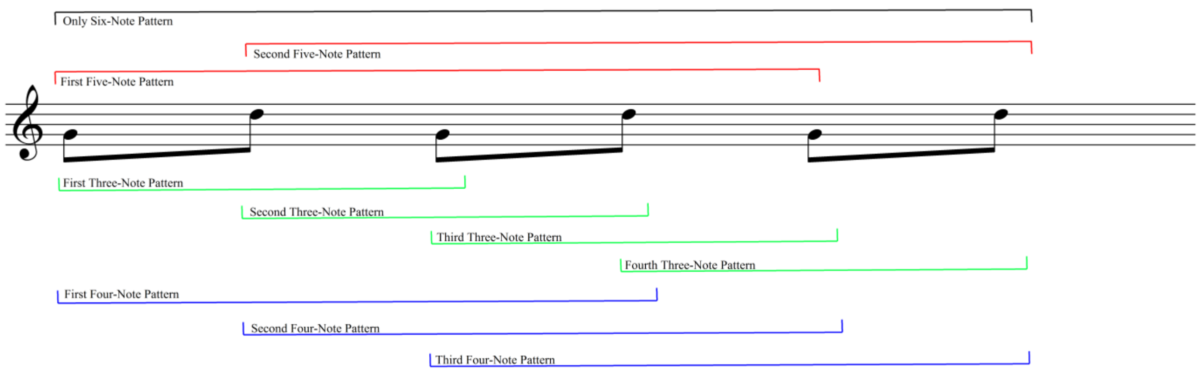

Example 5.4: Six alternating notes pattern generation before using the de-duplicating stored procedures

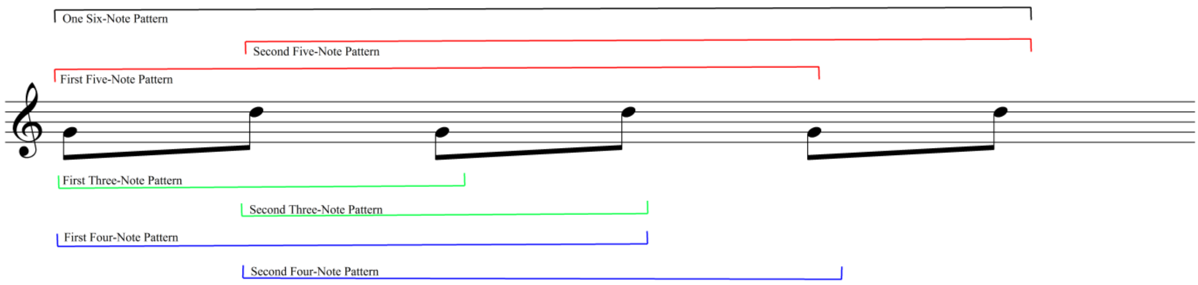

Example 5.5: Six alternating notes pattern generation using de-duplicating stored procedures

Unsurprisingly, this de-duplication process had a significant effect on the predominance of same-note PTs in the highest ranking positions. Table 5.2 is an amended version of Table 5.1 showing the data after the de-duplication process has taken place.

	Period									
	1		2		3		4		5	
PTD	Highest Ranked Same-Note PT	Nº in Top 20	Highest Ranked Same-Note PT	Nº in Top 20	Highest Ranked Same-Note PT	Nº in Top 20	Highest Ranked Same-Note PT	Nº in Top 20	Highest Ranked Same-Note PT	Nº in Top 20
A-PT	2	6	3	4	2	5	1	5	39	0
AD-PT	1	7	1	6	5	5	1	6	31	0
R-PT	5	2	5	2	3	2	6	2	10	1
RD-PT	4	4	3	2	4	2	4	2	10	1

Table 5.2: Predominance of same-note PTs in the 20 highest ranking positions after de-duplication

Table 5.2 clearly shows a reduction in the predominance of same-note PTs after the de-duplication process has taken place. Additionally, the highest ranking same-note PT is no longer always in the number one position. The difference in the predominance of same-note PTs between the last period and earlier periods is even more marked than before the de-duplication process, with the highest ranking position same-note A-PT and AD-PT being at positions 39 and 31 respectively.

5.4 Three-Note Pattern Types

The second factor influencing the highest ranking PTs is the number of notes in a pattern. Even after the de-duplication process, the highest ranking positions are dominated by three-note PTs. Table 5.3 shows the number of such PTs in the 20 highest ranking positions.

PTD	Period				
	1	2	3	4	5
A-PT	20	20	20	20	20
AD-PT	20	20	20	20	20
R-PT	15	16	17	19	17
RD-PT	14	15	15	19	16

Table 5.3: Number of three-note PTs in the 20 highest ranking positions

Table 5.3 clearly demonstrates the predominance of three-note PTs because for all the PTD/Period combinations, three-note PTs occupy the majority of the 20 highest ranking positions. In fact, all the 20 highest ranking position A-PTs and AD-PTs across all the periods consist of three notes.

The predominance of three-note PTs can be expected for two main reasons. Firstly, there are more three-note PTIs generated because there are more potential combinations of three notes within a composition than any other length of PTI (apart from two-note PTIs which are not being considered as part of this study). For example, a melodic line that has eleven notes will have only one eleven-note PTI compared to nine three-note PTIs. Secondly, there are fewer possible combinations of notes in a three-note PT. For example, a three-note pattern will have 12×12×12 possible combinations whereas a four-note pattern will have 12×12×12×12 possible combinations of notes (disregarding octave variations and enharmonic equivalents). The combination of there being far more potential three-note PTIs together with fewer possible combinations of notes within three-note PTs means that each potential three-note PT has a higher probability of occurring multiple times than longer-length PTs. This results in three-note PTs having a greater chance of prominence in the top ranking positions than their longer-length counterparts.

Table 5.3 also shows a difference in the predominance of three-note PTs in the 20 highest ranking positions between the absolute pitch value and the relative intervallic value PTDs. As noted, the absolute pitch value PTDs are completely dominated by three-note PTs in the 20 highest ranking positions for all periods, whereas the relative intervallic value PTDs are not (although there is an increase in the number of three-note PTs making the 20 highest ranking positions across the periods). As with the general predominance of three-note PTs, this is to be expected because there are fewer combinations of three-note relative intervallic value than absolute pitch value PTs. For example, with the A-PT and AD-PT there is a distinction between three consecutive Gs and three consecutive As, whereas with R-PT and RD-PT there is no such distinction.

5.5 Relative Frequency Rankings

To counteract the effect of shorter PTs having both a greater number of PTIs and a smaller number of potential combinations of notes, the ranking positions were recalculated to use a relative frequency figure rather than an actual frequency figure. The new calculation involved determining the frequency of each PT (using the number of PTIs multiplied by the number of movements in which the PT appears), then dividing the individual PT's frequency figure by the total number of PTIs for the appropriate PTL, multiplied by the total number of movements. Therefore, for each PTL and for each PT generated, the frequency calculation was

$$\frac{Op}{\sum_p Op} \times \frac{Mp}{\sum_p Mp}$$

where *Op* is the total number of times a PT occurs and *Mp* is the number of movements in which the PT was used.

For example, if a three-note PT occurs 22 times across 2 movements and the total number of all three-note PTIs is 1,000 across 20 movements, then the pattern type will have a relative frequency of 0.0022, i.e.

$$\frac{22}{1{,}000} \times \frac{2}{20} = 0.0022$$

Likewise, if a four note PT occurs 21 times across 2 movements and the total number of four-note PTs is 900 across 20 movements, then the pattern type will have a relative frequency of 0.0023, i.e.

$$\frac{21}{900} \times \frac{2}{20} = 0.0023$$

For these two examples, the four-note PT under the original ranking system would appear with a lower ranking position than the three-note PT, but under the new ranking system their positions would be reversed. Table 5.4 shows the number of three-note PTs within the 20 highest ranking positions using the new relative frequency algorithm.

	Period				
PTD	1	2	3	4	5
A-PT	20	20	20	20	20
AD-PT	20	20	20	20	20
R-PT	16	16	18	19	18
RD-PT	15	16	16	19	17

Table 5.4: Number of three-note PTs in the 20 highest ranking positions using the relative frequency algorithm

Table 5.4 shows that the recalculation of the ranking positions using relative frequencies has had no significant effect on the predominance of three-note PTs in the 20 highest ranking positions. For example, just like Table 5.3, both A-PTs and AD-PTs are completely dominated by three-note PTs in the 20 highest ranking positions for all periods. Likewise, the figures for R-PTs and RD-PTs are very similar to Table 5.3, with a maximum difference between comparable appearance figures of 2.

The fact that the using the new relative frequency calculation has had such a minimal effect is surprising. This is because the calculations earlier in this section on hypothetical three- and four-note PTs show how it is possible for a four-note PT to rise above a three-note PT in the ranking positions. To help investigate why there is not a great deal of difference in the 20 highest ranking positions when using the new as opposed to the original frequency algorithm, Table 5.5 shows the total number of PTIs for each PTL using A-PTs for each period.

	Period				
PTL	1	2	3	4	5
3	90,735	258,587	204,131	133,305	56,290
4	106,427	300,991	234,208	147,842	62,724
5	111,066	312,420	239,009	152,435	64,686
6	109,379	308,934	237,106	150,388	61,898
7	112,895	319,510	247,097	155,532	64,791
8	110,637	313,136	240,927	153,466	62,679
9	113,607	320,197	246,918	157,650	65,099
10	111,756	317,904	248,052	157,403	64,589
11	113,357	321,526	248,544	159,189	65,906

Table 5.5: Total number of PTIs for each PTL using A-PTs

Table 5.5 clearly shows that, for all periods, there are more four-note than three-note PTIs using A-PTs. This is counter-intuitive to the statement in Section 5.4 above that there should be more three-note than four-note PTIs. However, it must be remembered that the figures in Table 5.5 do not show all the possible PTIs. This is due, firstly, to all PTIs that begin or end with a rest being removed (see Section 4.6.6 above) and, secondly, because of the de-duplication process that removed PTIs that are duplicated when there is a repeated group of notes (see Section 5.3 above).

There is also a potential problem in using the relative frequency calculation. When calculating the absolute frequency of a pattern ($Op \times Mp$), the figure for the number of a PTIs is far greater than the figure for the number of movements in which a PT occurs (i.e., the left-hand figure in the calculation is *greater* than the right-hand figure). For example, the top ranking position three-note A-PT in period 1 has 762 instances across 60 movements (i.e., 762×60). This gives the number of PTIs a far greater weighting, especially when considering that there is a relatively small number of movements (442) compared to a large number of PTIs (7,834,928 for A-PTs across all periods). However, when using the relative frequency calculation ($\frac{Op}{\sum_p Op} \times \frac{Mp}{\sum_p Mp}$) the weightings between the figures for PTIs and movements are reversed, in that the movement figure will be far greater than the PTI figure (i.e. the left-hand figure in the calculation is now *smaller* than the right-hand figure). For example, using the previous highest ranking position three-note A-PT in period 1 example, the figure for the relative frequency of PTIs is 762 (the number of PTIs of this specific PT in period 1) divided by 90,735 (the total number of all three-note PTIs in period 1), resulting in a figure of 0.0084; whereas the figure for the frequency of the movements is 60 (the total number of movements in which the PT occurs within period 1) divided by 440 (the total number of possible movements in which the PT can occur), resulting in a figure of 0.1364. This means the figure for movements in this example has gone from being substantially smaller compared to the PTI figure (762 for the PTI figure against 60 for the movement figure) to being substantially larger (0.0084 for the PTI figure against 0.1364 for the movement figure).

Investigating the actual figures for some of the highest ranking position PTs may help to determine if this issue of the weighting between the number of PTIs and the number of movements in which PTs occur is a significant problem. Table 5.6 shows the figures for the ten highest three-note A-PTs in period 1, and Table 5.7 shows the same for four-note A-PTs.

Ranking Position	Nº of PTIs	Nº of Movements	Absolute Frequency	Relative Frequency
1	762	60	45,720	0.001145
2	586	75	43,950	0.001101
3	532	82	43,624	0.001093
4	710	59	41,890	0.001049
5	655	63	41,265	0.001034
6	658	61	40,138	0.001005
7	638	62	39,556	0.000991
8	668	58	38,744	0.000970
9	645	58	37,410	0.000937
10	526	69	36,294	0.000909

Table 5.6: Ten highest ranking position three-note A-PTs for period 1

Ranking Position	Nº of PTIs	Nº of Movements	Absolute Frequency	Relative Frequency
1	338	61	20,618	0.000440
2	337	60	20,220	0.000432
3	395	49	19,355	0.000413
4	285	65	18,525	0.000396
5	376	47	17,672	0.000377
6	354	47	16,638	0.000355
7	301	50	15,050	0.000321
8	299	48	14,352	0.000306
9	299	47	14,053	0.000300
10	330	37	12,210	0.000261

Table 5.7: Ten highest ranking position four-note A-PTs for period 1

Table 5.6 clearly shows that the number of movements that a PT occurs in has a direct impact on the ranking positions. For example, the PT ranked 3rd (highlighted in grey) in Table 5.6 occurs fewer times than that ranked 4th (also highlighted in grey), but the number of movements each PT occurs in is significantly greater for the PT ranked 3rd, giving it a greater absolute and relative frequency than that ranked 4th.

When comparing Table 5.6 to Table 5.7, it is noticeable that the figures for the number of PTIs and for the number of movements are, on the whole, greater for the three-note than for the four-note PTs. The figures for the number of PTIs and movements also result in values for both the absolute and relative frequencies being greater for the three-note than for the four-note PTs.

Because the comparison between three- and four-note PTs shows a clear difference in the number of PTIs as well as in the number of movements, and because there is little difference in the highest ranked PTs between using absolute and relative frequency figures , one must question whether ranking the PTs using the relative frequency algorithm is valid. Producing a relative frequency value to compare different PTLs is comparable to the concept of Relative Species Abundance, which concerns calculating how the population numbers of different species can be compared within different geographical locations. The complexities of the issues surrounding relative species abundance have resulted in a number of different theories for the calculation, none of which has become universally accepted ((Verberk, 2012) and (Volkov, et al., 2003)). Consequently, the idea of trying to determine a single meaningful figure for the distribution of an individual PT is beyond the scope of this study and will not be investigated any further. Instead, the present study will revert to just using the absolute frequency ranking positions within each PTL when comparing the ranking positions of PTs.

5.6 Pattern Type Shape Property

Chapter 4, Section 4.6.7, provides an explanation for how the PT pitch shape property was calculated (it shows the overall direction of the movement between notes). Same-note PTs produce a shape property of 'S' and this shape has been shown to have an effect on the ranking positions, as shown in Section 5.2. However, there are many other PT shape properties identified that need to be investigated in terms of their effect on the ranking positions. Table 5.8 shows the number of each PT shape property within the 20 highest ranking position three-note PTs.

PTD	Shape Property[11]	Period				
		1	2	3	4	5
A-PT	D	12	12	9	9	6
	DU	2	3	1	1	4
	S	6	4	5	5	-
	U	-	1	5	4	8
	UD	-	-	-	1	2
AD-PT	D	11	11	10	7	5
	DU	1	2	1	1	2
	S	7	6	5	6	-
	U	1	1	4	5	9
	UD	-	-	-	1	4
R-PT	D	4	3	3	3	4
	DS	2	2	2	2	1
	DU	3	3	2	2	3
	S	1	1	1	1	1
	SD	2	2	2	2	1
	SU	2	1	2	2	2
	U	3	3	3	4	4
	UD	2	4	3	2	3
	US	1	1	2	2	1
RD-PT	D	4	3	3	3	4
	DS	2	2	2	2	1
	DU	2	3	3	2	3
	S	1	1	1	1	1
	SD	2	2	2	2	1
	SU	2	1	2	2	1
	U	3	3	3	4	4
	UD	3	5	3	3	4
	US	1	-	1	1	1

Table 5.8: Number of each shape property for the 20 highest ranking position three-note PTs

[11] An explanation of the shape properties is provided in Section 4.4.4.

Table 5.8 confirms that, after de-duplication, same-note PTs are no longer dominant (i.e., most frequent) within the 20 highest ranking positions for three-note PTs. In fact, for A-PTs and AD-PTs, the dominant (i.e., most frequent) shape property is now a unidirectional downward progression (shown in Example 5.6) across the first four periods. However, for R-PTs and RD-PTs, there is no particularly dominant shape property across any of the periods.

Example 5.6: Top-ranked three-note pattern type with a shape property of 'D' in period 1 within the dataset

There are also a higher number of different shape properties for R-PTs and RD-PTs (nine in total) than for A-PTs and AD-PTs (five in total) in the 20 highest ranking position three-note PTs. This difference between the absolute pitch value and the relative intervallic value PTDs can be explained by the fact that the relative intervallic PTDs encompass twelve different absolute pitch value PTDs. For example, a pattern of three notes that rises through the first three notes of a major scale can appear as a number of separate PTs under the absolute pitch value PTDs because each major scale will count as a different PT. However, the same pattern will appear as only one PT under the relative intervallic value PTDs because the intervals are the same regardless of the key.

Table 5.8 shows the shape property for three-note PTs but there are a limited number of shape properties that can be generated for three-notes. This is because there can only be a maximum of two changes in the direction between pitches for three-note PTs. However, a longer-length PT will have a greater number of possible changes in direction: for example, seven and eleven-note PTs can have up to six and ten changes in direction respectively. Table 5.9 shows the shape properties for a longer-length PT (seven notes) in the 20 highest ranking positions.

PTD	Shape Property	Period				
		1	2	3	4	5
A-PT	D	4	10	3	-	-
	DS	1	-	-	1	-
	DU	1	-	-	-	-
	DUD	-	-	-	-	2
	DUDU	-	-	-	-	1
	DUDUDU	-	-	-	-	3
	S	9	10	11	18	10
	SD	3	-	-	-	-
	SDS	3	-	-	-	-
	SU	-	-	1	-	-
	SUS	1	-	-	-	-
	U	-	-	5	-	3
	UDU	-	-	-	-	1
	UDUD	-	-	-	-	1
	UDUDUD	-	-	-	1	-
	US	-	-	-	1	-
AD-PT	D	-	10	5	-	-
	DUD	-	-	-	-	5
	DUDU	-	-	-	-	2
	DUDUDU	-	-	-	6	2
	S	10	10	11	13	6
	SD	3	-	-	-	-
	SDS	12	-	-	1	-
	SUS	1	-	-	-	-
	U	-	-	7	1	3
	UDU	-	-	-	-	2
	UDUD	-	-	-	-	4
	UDUDUD	-	-	-	7	-
R-PT	D	4	7	5	-	1
	DS	2	-	1	1	1
	DU	1	1	-	-	-
	DUDU	-	-	-	-	2
	DUDUD	-	-	-	-	1
	DUDUDU	-	1	-	3	2
	S	1	2	2	6	7
	SD	2	1	2	2	-
	SDS	6	1	3	3	-
	SU	1	1	2	1	1
	SUS	2	2	2	-	-
	U	-	3	2	-	4
	UDUD	-	-	-	-	2
	UDUDUD	-	-	-	3	1
	US	1	1	1	1	-
RD-PT	D	-	7	7	2	-
	DS	2	-	-	1	-
	DUD	-	-	1	-	-

PTD	Shape Property	Period				
		1	2	3	4	5
	DUDU	-	-	-	-	4
	DUDUD	-	-	-	-	1
	DUDUDU	-	-	-	4	2
	S	1	1	1	3	2
	SD	2	1	2	1	-
	SDS	8	4	3	2	-
	SU	1	1	2	1	-
	SUS	6	2	1	1	-
	U	-	3	4	1	4
	UDUD	-	-	-	-	4
	UDUDU	-	-	-	-	2
	UDUDUD	-	-	-	4	2
	US	1	1	-	-	-

Table 5.9: Number of each shape property for the 20 highest ranking position seven-note PTs

Table 5.9 shows that for seven-note PTs, the same-note shape property for A-PTs and AD-PTS becomes more prevalent, with a reduction in the unidirectional downward progression dominance compared with three-note PTs (Table 5.8). However, the other findings of the three- and seven-note PT shape properties are similar. For example, there is a greater diversity of shape properties for R-PTs and RD-PTs than for A-PTs and AD-PTs, and there are no dominant shape properties for R-PTs and RD-PTs.

When looking at the data in Table 5.8 and Table 5.9 for all four PTDs, the unidirectional shape properties of D, S and U as a whole seem to dominate. Table 5.9 gives the frequency of PTs with a unidirectional shape property within the 20 highest ranking positions for three-, seven- and eleven-note PTs within the dataset.

Table 5.10 shows that for three-note A-PTs and AD-PTS there is a preponderance of unidirectional shape properties. However, the unidirectional downward progression three-note PTs are clearly declining in number over the periods. Conversely, the unidirectional upward progression PTs are clearly increasing in number across the periods. For the three-note R-PTs and RD-PTs a marked decrease in either the unidirectional downward or upward progressions is not evident. When looking at the seven-note PTs, A-PTs and AD-PTs show a preponderance of same-note PTs, whereas for R-PTs and RD-PTs no unidirectional movement type is dominant. Finally, the eleven-note PTs are dominated by the same-note shape property across all four PTDs, although for R-PTs and RD-PTs the figures are small considering they are out of a total of 20 ranking positions.

			Period				
PTL	PD	Shape	1	2	3	4	5
3	A-PT	D	12	12	9	9	6
		S	6	4	5	5	-
		U	-	1	5	4	8
	AD-PT	D	11	11	10	7	5
		S	7	6	5	6	-
		U	1	1	4	5	9
	R-PT	D	4	3	3	3	4
		S	1	1	1	1	1
		U	3	3	3	4	4
	RD-PT	D	4	3	3	3	4
		S	1	1	1	1	1
		U	3	3	3	4	4
7	A-PT	D	4	10	3	-	-
		S	9	10	11	18	10
		U	-	-	5	-	3
	AD-PT	D	-	10	5	-	-
		S	10	10	11	13	6
		U	-	-	7	1	3
	R-PT	D	4	7	5	-	1
		S	1	2	2	6	7
		U	-	3	2	-	4
	RD-PT	D	-	7	7	2	-
		S	1	1	1	3	2
		U	-	3	4	1	4
11	A-PT	D	-	1	1	-	-
		S	9	10	12	17	14
		U	-	-	-	-	4
	AD-PT	D	-	-	-	-	-
		S	10	10	11	13	12
		U	-	-	-	-	2
	R-PT	D	-	1	-	-	-
		S	1	1	1	4	6
		U	-	-	-	-	-
	RD-PT	D	-	-	-	-	-
		S	2	1	1	2	3
		U	-	-	-	-	-

Table 5.10: Number of each unidirectional shape property for the 20 highest ranking position three-, seven- and eleven-note PTs

The fact that same-note PTs seem to become more predominant when increasing the PTL suggests that the de-duplication process has a greater effect on shorter-length patterns than on their longer-length counterparts. This accords with the de-duplication methodology, in that seven consecutive notes on the same pitch would originally have generated five three-note patterns and only one seven-note pattern, whereas after the de-duplication process the same seven-note pattern would still generate one seven-note pattern but only one three-note pattern (see Section 5.3).

5.7 Summary

The initial observations have shown that there is a predominance of three-note, same-note PTs in the 20 highest ranking positions. However, whilst the predominance of three-note PTs is spread across all the periods, it is not as marked in the final period. After the de-duplication process, the number of three-note same-note PTs in the 20 highest ranking positions is reduced. However, Section 5.6 shows that this reduction is not as marked when looking at the longer-length A-PTs and AD-PTs.

To compensate for the dominance of three-note PTs, an attempt was made to relativize the different PTLs by creating a single value for the relative frequency of each PT compared to other length PTs. Unfortunately, this illustrated the complexities involved in creating a meaningful figure that would show the distribution of PTIs across all PTLs. As such, it was decided to compare only PTs of the same length.

Finally, the PT shape property was investigated. This found that for three-note A-PTs and AD-PTs there is a decrease in the frequency of the unidirectional downward progression and an increase in the unidirectional upward progression shape properties across the periods. However, this is not the case for the seven- and eleven-note PTs, where for the majority of PTD/period combinations the same-note PT is predominant.

Given that the initial observations highlight that there are some changes taking place across the periods, particularly in the PT shape property, it must be asked whether these changes are the result of the evolutionary processes of selection, replication and variation, and subsequently do they provide evidence to support the concept of memes in music.

Part III – Analysis

6 Chapter 6: Evolutionary Processes

6.1 Introduction

The pattern with the highest ranking position in period 1 (i.e., the pattern with the most occurrences in the most movements) is three consecutive Gs in the same octave, when using the absolute pitch values of the notes. In fact, all the periods have in their highest ranking position a pattern that consists of three consecutive notes the same, whether looking at absolute pitch values or relative intervallic values, whether including the durational value of the notes or not. However, the three Gs from period 1 do not remain in the highest ranking position across all five periods. In periods 2 and 3 they do, but in period 4 the pattern drops through the ranking positions to 8th place and in period 5 drops further to 16th place. Therefore, the pattern of three Gs becomes less frequent through the periods in relation to other patterns, and consequently it can be argued that it is being selected and replicated fewer times across the periods. Additionally, there is a different three-note pattern consisting of consecutive notes the same (three Ds), which could be considered a variant of three Gs, appearing in the highest ranking position for periods 4 and 5. But does this movement through the ranking positions of the three Gs, together with a possible variant pattern of three Ds replacing the three Gs in the highest ranking position, make the three Gs a meme in music?

As argued in Chapter 2, for memes to exist they need to display the evolutionary processes of selection, replication and variation, as well as exhibit the replicator properties of longevity, fecundity and copying-fidelity. If it is accepted that a pattern in music is a reasonable unit of information, and that a meme can be a unit of information, then patterns in music should show the evolutionary processes as well as exhibit the replicator properties. The methodology outlined in Part II results in a series of ranking positions for patterns in music based upon the frequency of their occurrences within the dataset across the periods. It is these ranking positions that are explored in this chapter to investigate whether there is any evidence for patterns in music displaying the evolutionary processes of selection, replication and variation.

Selection and replication are investigated together when analysing the ranking positions, since selection is a moderating force on replication. It is argued that if some patterns are consistently at the top of the ranking positions, they are being used more frequently than those patterns at the bottom of the ranking positions, and consequently the former are being selected and replicated more widely than the later. Additionally, if patterns are consistently progressing upwards through the ranking positions, those patterns are showing signs of being selected and replicated an increasing number of times across the periods. Likewise, the opposite is true for patterns consistently progressing downwards through the ranking positions. Some of the patterns that show either a consistent progression upwards or downwards through the ranking positions are then sampled to investigate the various scenarios in which the patterns appear. This further investigation shows that the patterns could appear in a number of different structural places within a quartet (in expositions, in codas, in trio sections, as part of cadences, etc.), in any of the instrumental parts, and in a number of different scenarios (such as part of the melodic interest, part of an accompaniment figure, part of a bass line, part of a sequence, part of a parallel passage between instruments, etc.).

Evidence for variation is explored using the similarity algorithm explained in Section 3.3 above. Firstly, an investigation is made into the effectiveness of the similarity algorithm used. This investigation highlights the difficult nature of defining similarity, and shows both the positive and negative aspects of the algorithm in that it provides both possible and improbable connections between patterns. The next stage involves taking patterns that first appeared in period 3, seeing if those patterns progressed upwards or downwards through the subsequent periods, and finding possible connections with patterns from the previous two periods.

The investigations find that there are some new patterns in period 3 that, through the ranking positions, provide evidence of selection and replication. However, the numbers of patterns involved is extremely small compared to the total number of patterns within the dataset. Likewise, some evidence is found for connections between the new patterns in period 3 and some patterns in the earlier periods, but the similarity algorithm is not completely effective.

6.2 Selection and Replication

A composer's works are shaped by his or hers place in history; the composer creates works that accord with a progression from the past, through the present, and lead to the future. In order to achieve this, a composer is (either intentionally or unintentionally) selecting and replicating elements from the past to fit into a present context. Part of this process could be the selection and replication of patterns in music that the composer has heard and believes work as part of his or her own compositional style or as part of an individual composition. In other words, the composer is selecting and replicating material perceived to be relevant to their own work.

When PTs are selected and replicated by composers, some should gain more prominence in the ranking positions across the periods. This is because those PTs that are used by composers more frequently across time will, self-evidently, gain a higher ranking position across the periods than those PTs that are not. Additionally, it should also be possible to see some PTs disappearing from the ranking positions in the later periods if they are no longer being selected and replicated. There should also be PTs that appear in the ranking positions in only one or two periods, because new PTs are not necessarily selected and replicated. However, the movement of the PTs within the ranking positions across the periods is not by itself conclusive evidence for selection and replication.

Jan points out that, '[t]hose memes that have the greatest perceptual-cognitive salience will tend to be the most widely propagated in the meme pool, to the disadvantage of those less salient forms' (2007, p. 229). However, this raises the question of what gives a meme 'perceptual-cognitive salience'? For example, the PT in Example 5.1 has arguably been selected and replicated widely (it is a top-ranked PT), but does the PT have any defining characteristics that give it 'perceptual-cognitive salience'? If the three Ds were part of an accompaniment figure, then would the pattern type have any more or less 'perceptual-cognitive salience' than if they were at the start of both the first and second subject of a sonata-form movement?

It could be argued that because the PT of three Ds is at the top of the ranking positions then it automatically has 'perceptual-cognitive salience', due to the fact that it has been widely selected and replicated across compositions. However, the PT of three Ds could simply be a useful tool in the composer's palette, or it could be that it is part of the basic fundamentals of composition and will therefore be found widely across compositions, regardless of their chronology (i.e., Dennett's 'Good Trick' (1995, pp. 77-78)).

Additionally, there is the issue of whether a composer had knowledge of a particular PT before incorporating it into a composition. For example, has composer X used a PT with or without any knowledge that composer Y has already used it? However, just because composer X might not have any knowledge of composer Y does not necessarily mean that there is no link between the two composers. For example, it could be that both composer X and Y heard the PT from a composition by composer Z. Unfortunately, stemma-like investigations into the connections between different composers are beyond the scope of the present research.

What can be investigated is the movement between ranking positions of the PTs across the periods. This movement can then be used to show if any PTs are becoming more, or less, dominant across time, helping to provide evidence for possible selection and replication.

Another aspect that can be investigated is the type of scenarios in which the PTs appear. For example, it could be that those PTs that progress upwards through the ranking positions are used in similar scenarios, whereas those that progress downwards are used in differing scenarios. If this is the case it would show that the function of the PTs within the music has an impact on selection and replication.

6.2.1 Exploring Selection and Replication within the Dataset

To explore each PTs progression through the ranking positions, a calculation was made as to how many times each PT moved upwards, downwards, or stayed in the same ranking position between consecutive periods, either with or without appearing in the appropriate first or last period according to their progression direction, as shown in Table 6.1.

	Consistent Upwards Progression		Consistent Downwards Progression	
PTD	With Appearance in All Periods	With Appearance in Periods 2 to 5 Only	With Appearance in All Periods	With Appearance in Periods 1 to 4 Only
A-PT	44	5,188	30	46
AD-PT	20	2,911	17	28
R-PT	53	5,180	54	36
RD-PT	36	3,333	73	43

Table 6.1: Number of PTs consistently progressing upwards or downwards through the ranking positions across the periods

The figures in Table 6.1 show that for all four PTDs there are some PTs that consistently progress either upwards or downwards through the ranking positions across the periods. However, these figures are extremely small when compared to the total number of PTs for each PTD. For example, there are over four million A-PTs, making the percentage figure for the number of PTs that consistently progress upwards through the ranking positions across the periods (44) equal to 0.001%. Even the highest number within Table 6.1 of PTs (5,188) produces a small percentage figure (0.124%) when compared to the total number of PTs for the appropriate PTD (greater than four million).

Table 6.1 also shows that for all the PTDs there are a greater number of PTs that do not appear in period 1 and then consistently progress upwards through the ranking positions across the periods than the number of PTs doing the opposite, i.e., progressing downwards through the ranking positions and not appearing in period 5. This would suggest that, compared to the number of PTs that consistently progress upwards through the ranking positions across all the periods, there are a larger number of PTs that were not used by composers in period 1, but were then used in period 2, and then gained more popularity during subsequent periods. However, the reverse is not true, in that there are not greater numbers of PTs that completely fall out of use in period 5 having consistently progressed downwards through the ranking positions during the previous four periods.

The dataset also shows that there are PTs that are not used in all five periods. Table 6.2 shows the number of PTs according to the number of periods in which each PT appears.

	Number of Periods in which each PT Appears (with % of all PTs)									
PTD	In Only One out of Five		In Only Two out of Five		In Only Three out of Five		In Only Four out of Five		In All Five	
A-PT	3,906,786	93.5%	170,300	4.1%	55,570	1.3%	28,433	0.7%	16,162	0.4%
AD-PT	4,952,734	96.6%	113,375	2.2%	35,280	0.7%	16,566	0.3%	8,406	0.2%
R-PT	2,802,562	91.9%	152,376	5.0%	51,227	1.7%	25,925	0.9%	16,590	0.5%
RD-PT	3,967,525	96.0%	104,786	2.5%	33,577	0.8%	15,937	0.4%	9,045	0.2%

Table 6.2: Number of PTs appearing in one or more periods

Table 6.2 shows that, for all PTDs, the overwhelming majority of PTs appear in only one out of the five periods. For example, the A-PTs show that 93.5% of PTs appear in only one period and only 0.4% appear in all five periods. Additionally, Table 6.2 also shows that, for all four PTDs, there is a inverse correlation between the number of PTs appearing in either one, two, three, four or five periods, and the number of periods in which PTs appear; i.e., there are fewer PTs appearing in all five periods than in four out of five periods, and there are fewer appearing in four out of five periods than in three out of five periods, etc.

Section 5.2 showed how three-note PTs dominated the ranking tables when comparing frequencies of different PTLs. This implies that length has an effect on the frequencies with which individual PTs are used by composers. Does this affect how different PTLs behave according to the number of periods they appear in and how they progress through the ranking positions across the periods? That is, do all PTLs have a relatively small number of PTs consistently progressing upward or downward through the ranking positions, do all PTLs have an overwhelming majority of PTs appearing in just one period, and do all PTLs have a relatively small number of PTs appearing in all five periods, across all four PTDs? Table 6.3 is similar to Table 6.1 in that it shows the number of PTs consistently progressing upwards or downwards through the ranking positions across periods, but here the data is organized according to PTL.

		Consistent Upwards Progression		Consistent Downwards Progression	
PTD	PTL	With Appearance in All Periods	With Appearance in Periods 2 to 5 Only	With Appearance in All Periods	With Appearance in Periods 1 to 4 Only
A-PT	3	18	824	7	2
	4	11	1,728	5	4
	5	8	1,247	14	1
	6	4	600	3	4
	7	2	364	1	7
	8	-	173	-	5
	9	1	120	-	10
	10	-	80	-	5
	11	-	53	-	8
AD-PT	3	12	1,032	6	2
	4	5	876	7	8
	5	2	477	4	7
	6	1	242	-	3
	7	-	142	-	5
	8	-	63	-	2
	9	-	33	-	1
	10	-	29	-	-
	11	-	17	-	-
R-PT	3	14	5	20	2
	4	14	804	9	3
	5	9	1,450	9	3
	6	3	1,151	5	5
	7	3	652	3	4
	8	1	435	1	3
	9	2	313	3	6
	10	4	217	1	5
	11	3	153	3	5
RD-PT	3	11	617	12	7
	4	8	904	14	7
	5	5	604	13	4
	6	5	401	11	6
	7	1	300	8	4
	8	1	204	6	4
	9	2	139	4	5
	10	2	100	1	2
	11	1	64	4	4

Table 6.3: Number of PTs consistently progressing upwards or downwards through the ranking positions across the periods organised by PTL

While the majority of the figures in Table 6.3 for PTs progressing upwards or downwards are too small to be meaningful, on the whole they accord with the findings from Table 6.1, in that they are relatively small compared with the overall number of PTs. Similarly, the column with the largest values are for those PTs that do not appear in period 1 but then consistently progress upwards through the ranking positions across periods from period 2 onwards.

There is, however, an interesting quirk in the figures for PTs that do not appear in period 1 but then consistently progress upwards through the ranking positions across periods. It is not always the three-note PTs that have the largest number of PTs following this trend. For example, the four-note A-PTs and RD-PTs, and the five-note R-PTs are the most frequent PTs following this upward trend. In fact, only five three-note R-PTs do not appear in period 1 and then progress upwards through the ranking positions across periods compared with 1,450 five-note R-PTs.

Dividing Table 6.1 according to PTLs revealed some significant differences between the different lengths. Further significant results are produced when dividing Table 6.2 according to PTL, as shown in Table 6.4 below.

Table 6.4 shows that not all PTLs produce the same results as the data as a whole (see Table 6.2). There is a distinction between how the shorter and longer-length PTs behave within Table 6.4. For example, when looking at the data for PTs as a whole, there is an inverse correlation between the number of PTs appearing in only one, two, three, four or five periods, and the number of periods involved; i.e., as the number of periods increases, the number of PTs decreases. However, this is not always the case when taking into account specific PTLs. For example, the three-note A-PTs, R-PTs and RD-PTs, and the four-note R-PTs do not follow this correlation.

Also, the percentage of PTs that appear in all five periods compared to all PTs within that PTD changes significantly according to PTL. For example, the percentage figure for three-note A-PTs appearing in all five periods is 16.9% compared to the equivalent percentage figure for eleven-note PTs of 0.0%. This difference between PTLs shows a correlation between the PTL and the percentage figure of PTs appearing in all five periods, with the longer PTLs having a smaller percentage figure. The difference between the shorter and longer PTLs is true for all four PTDs, although for the AD-PTs and RD-PTs the difference in the percentage figures between thee- and eleven-note PTLs is smaller (3.3 and 6.4 percentage points respectively) than for A-PTs and R-PTs (16.9 and 34.5 percentage points respectively) in percentage-point terms.

Another point to make regarding Table 6.4 is that it is not always the three-note PTLs that have the greatest absolute frequency of PTs appearing in all five periods. For example, the four-note A-PTs, R-PTs and RD-PTs all have a greater absolute frequency figure for PTs appearing in all five periods than their three-note counterparts. However, the three-note PTs have the greatest percentage figure of PTs appearing in all five periods for all four PTDs.

		Number of Periods in which each PT Appears (with % of all PTs)									
		In Only One out of Five		In Only Two out of Five		In Only Three out of Five		In Only Four out of Five		In All Five	
PTD	PTL										
A-PT	3	8,183	44.9%	3,029	16.6%	1,961	10.8%	1,956	10.7%	3,086	16.9%
	4	59,933	59.2%	17,613	17.4%	10,227	10.1%	7,685	7.6%	5,747	5.7%
	5	201,559	75.6%	38,160	14.3%	15,514	5.8%	7,816	2.9%	3,559	1.3%
	6	379,234	87.4%	37,600	8.7%	10,753	2.5%	4,361	1.0%	1,774	0.4%
	7	525,268	93.1%	28,436	5.0%	6,948	1.2%	2,769	0.5%	1,003	0.2%
	8	606,717	96.2%	17,888	2.8%	4,092	0.6%	1,570	0.2%	460	0.1%
	9	675,490	97.6%	12,723	1.8%	2,817	0.4%	1,041	0.2%	265	0.0%
	10	709,142	98.4%	8,532	1.2%	1,833	0.3%	705	0.1%	155	0.0%
	11	741,260	98.9%	6,319	0.8%	1,425	0.2%	530	0.1%	113	0.0%
AD-PT	3	78,841	72.5%	14,669	13.5%	7,068	6.5%	4,519	4.2%	3,607	3.3%
	4	290,362	86.6%	28,237	8.4%	9,725	2.9%	4,694	1.4%	2,390	0.7%
	5	462,302	93.2%	23,081	4.7%	6,510	1.3%	2,725	0.5%	1,260	0.3%
	6	564,334	96.5%	14,832	2.5%	3,706	0.6%	1,597	0.3%	623	0.1%
	7	645,970	97.7%	11,079	1.7%	2,896	0.4%	1,226	0.2%	295	0.0%
	8	677,574	98.5%	7,551	1.1%	1,958	0.3%	754	0.1%	113	0.0%
	9	724,403	98.9%	6,041	0.8%	1,474	0.2%	483	0.1%	56	0.0%
	10	742,173	99.2%	4,368	0.6%	1,077	0.1%	328	0.0%	30	0.0%
	11	766,775	99.4%	3,517	0.5%	866	0.1%	240	0.0%	32	0.0%
R-PT	3	544	28.7%	289	15.2%	204	10.7%	207	10.9%	654	34.5%
	4	8,266	43.1%	3,151	16.4%	2,209	11.5%	2,168	11.3%	3,405	17.7%
	5	54,319	59.5%	16,286	17.8%	9,251	10.1%	6,513	7.1%	4,938	5.4%
	6	174,438	76.1%	32,468	14.2%	12,699	5.5%	6,362	2.8%	3,108	1.4%
	7	327,473	86.8%	33,268	8.8%	10,439	2.8%	4,385	1.2%	1,892	0.5%
	8	448,209	92.6%	25,427	5.3%	6,569	1.4%	2,494	0.5%	1,084	0.2%
	9	541,846	95.4%	18,941	3.3%	4,553	0.8%	1,719	0.3%	768	0.1%
	10	600,407	97.2%	12,914	2.1%	3,063	0.5%	1,185	0.2%	441	0.1%
	11	647,060	98.0%	9,632	1.5%	2,240	0.3%	892	0.1%	300	0.0%
RD-PT	3	22,458	63.8%	5,475	15.6%	2,894	8.2%	2,118	6.0%	2,256	6.4%
	4	143,217	80.3%	20,018	11.2%	7,955	4.5%	4,413	2.5%	2,854	1.6%
	5	302,391	89.7%	22,282	6.6%	7,400	2.2%	3,326	1.0%	1,626	0.5%
	6	427,180	94.6%	16,953	3.8%	4,754	1.1%	1,940	0.4%	803	0.2%
	7	522,087	96.5%	13,166	2.4%	3,518	0.7%	1,364	0.3%	621	0.1%
	8	573,709	97.8%	8,970	1.5%	2,370	0.4%	974	0.2%	383	0.1%
	9	628,712	98.4%	7,478	1.2%	1,968	0.3%	796	0.1%	250	0.0%
	10	657,849	98.8%	5,756	0.9%	1,497	0.2%	568	0.1%	143	0.0%
	11	689,922	99.1%	4,688	0.7%	1,221	0.2%	438	0.1%	109	0.0%

Table 6.4: Number of PTs appearing in one or more periods organised by PTL

It can be argued that both the progression tables (Table 6.1 and Table 6.3) and the period appearance tables (Table 6.2 and Table 6.4) show evidence for selection and replication. Table 6.1 and Table 6.3 show that some PTs (albeit a relatively small number compared to the total number of PTs) either show a consistently upwards or downwards progression through the ranking positions across the periods. This would be consonant with the hypothesis of PTs being selected and replicated, with those PTs progressing upwards through the ranking positions being selected and replicated more widely than other PTs across the periods, and those PTs progressing downwards through the ranking positions being replicated and selected fewer times than other PTs across the periods. However, the small numbers involved compared to the total number of PTs does call into question the significance of these PTs as evidence of selection and replication.

The small number of PTs consistently progressing upwards or downwards through the ranking positions could be a result of the methodology. As explained in Section 4.6.6, all possible patterns from the dataset were generated to accord with the working definition of a meme in music. This will have generated patterns that cross the boundaries of what a composer might consider a pattern in music (such as phrasing, bar-lines, rests, etc.). Therefore, the methodology may have generated a significant number of patterns that, musically, may not be valid. For example, a pattern that consists of the last three notes of the Minuet section and the first three notes of the Trio section from a Minuet and Trio movement will be treated as a six-note pattern by the pattern generation algorithm. Consequently, these *extra* patterns generated will affect the ratio between the number of PTs that continuously progress upwards of downwards through the ranking positions and the total number of PTs.

Table 6.2 and Table 6.4 also show evidence for selection and replication. The fact that there are PTs that only appear in one period is consistent with the argument that these PTs are being created, but are not subsequently being selected and replicated. Again, there is a problem with the numbers of PTs because the overwhelming majority only appear in one period with relatively few appearing in other periods. This issue with the small number of PTs appearing in more than one period could also be a result of the methodology generating patterns that may potentially not be considered valid.

The fact that each PTL can produce differing results can also be interpreted as evidence to support the hypothesis of selection and replication. The data shows that shorter PTs are being used more frequently than their longer counterparts. This could be a result of shorter PTs being more susceptible to selection and replication owing to their being more easily recalled and amalgamated into new compositions. Nevertheless, the numbers of PTs involved is again relatively small compared to the total number of PTs.

6.2.2 Exploring Selection and Replication in the Music

This section investigates the placement within the music of some of the PTs in order to ascertain whether there is any link between the ranking positions and the context in which PTs appear. It would, perhaps, be expected that PTs at the top of the ranking positions would appear in similar contexts, whereas PTs at the bottom of the ranking positions could be expected to occur in a variety contexts. Therefore, a small number of random examples of PTIs are explored for their context within the music.

Before investigating the contexts of the PTIs, it should be noted that looking at the data *en masse* will not always give an accurate picture of the behaviour of PTs within the ranking positions across the periods when looking for evidence of selection and replication. For example, a PT can progress upwards (or downwards) through the ranking positions across the periods by moving either upwards (or downwards) just one position between each pair of consecutive periods. This means that the PT is actually staying in roughly the same position across the periods and is arguably not showing as much evidence for selection and replication as a one that moves from the bottom to the top (or vice versa) of the ranking positions across the periods. This is because the number of PTIs increases (or decreases) more across the periods for PTs that progress through a large number of ranking positions compared to those that stay in roughly the same position.

It is also possible for a PT to show an upward (or downward) progression through the ranking positions across the periods but to still remain towards the bottom of the ranking positions, meaning that the PT is remaining relatively obscure. If a PT remains at the bottom of the ranking positions, then it is not being selected and replicated as widely as one that moves from the bottom to the top of the ranking positions, because the ranking positions are based, of course, on the number of PTIs.

In order to overcome these problems of placement within the ranking positions when investigating PTs within the music, it is necessary to filter out any PTs that do not progress very far, or which remain close to the bottom of the ranking positions. Table 6.5 shows the total number of PTs consistently progressing either upwards or downwards through the ranking positions across the periods, how many of them either start or end their progression within the 50 highest ranking positions depending on the direction of progression, how many of them progress through at least 100 ranking positions between the first and last periods, and how many fit both these criteria.

Progression Direction	PTD	Total Nº of PTs	Nº Starting or Ending in Top 50	Nº Progressing through at least 100 Ranking Positions	Nº Fitting Both Criteria
Upwards	A-PT	44	13	37	6
	AD-PT	20	8	16	4
	R-PT	53	16	40	9
	RD-PT	36	17	30	12
Downwards	A-PT	30	18	24	15
	AD-PT	17	10	12	6
	R-PT	54	24	33	8
	RD-PT	73	50	43	24

Table 6.5: Number of PTs consistently progressing upwards or downwards through the ranking positions across the periods

Table 6.5 shows that the majority of PTs investigated in the table progress a significant distance through the ranking positions across the periods, and that there are a substantial number of PTs that start or end (according to their progression direction) in the 50 highest ranking positions. This indicates that most of the PTs do not progress through only a small number of ranking positions, nor do they remain towards the bottom of the ranking positions.

Table 6.6 displays the PT moving through the greatest number of ranking positions across the periods for each progression direction and PTD, whereas Table 6.7 displays the PT that has the highest ranking position in either the first or last period depending on the progression direction. PTs that progress upwards are highlighted in grey and PTs that progress downwards are not highlighted.

There are a couple of major differences between Table 6.6 and Table 6.7 in terms of the PTL and directional movement of the PTs involved. In Table 6.6, the PTs are on the whole shorter than those in Table 6.7. For example, all the PTs in Table 6.6 are either three or four notes long, whereas in Table 6.7 the PTs range from three to eleven notes long. Likewise, the PTs in Table 6.6, compared to Table 6.7, have on the whole more disjunctive movement. For example, all the PTs in Table 6.7 have a movement of either one or two semitones or they stay on the same note, whereas those in Table 6.6 tend to make use of some larger intervals. However, when looking at the durational aspect of the PTs, both tables show that seven out of the eight PTs that include duration (i.e., AD-PTs and RD-PTs) use the same length note throughout.

PTD	PT[12]	Shape Property	PTL	Ranking Position within Period				
				1	2	3	4	5
A-PT	A/D0A-1/D0G	D	3	629	522	487	209	92
AD-PT	A:1/D0G+1:1/S0G+1:1	DS	3	449	394	237	228	77
R-PT	0/U1/U3/D2	UD	4	805	667	592	439	173
RD-PT	0:1/D4:1/U3:1	DU	3	550	527	359	167	113
A-PT	G/D0E-1/D0D	D	3	124	359	392	415	421
AD-PT	G:1/S0G:1/S0G:2	S	3	137	190	226	352	369
R-PT	0/D5/D3/D4	D	4	150	174	232	392	450
RD-PT	0:1/D12:1/U5:1	DU	3	86	182	218	320	435

Table 6.6: PTs that consistently progress upwards (highlighted in grey) or downwards through the ranking positions with the largest difference in ranking positions between the first and last periods

PTD	PT[12]	Shape Property	PTL	Ranking Position within Period				
				1	2	3	4	5
A-PT	D/U0E/U0F/U0G	U	4	82	44	15	10	1
AD-PT	D:1/U0E:1/U0F:1/U0G:1	U	4	79	46	19	11	1
R-PT	0/U2/U1	U	3	8	7	4	3	2
RD-PT	0:1/D1:1/U1:1/D1:1/U1:1 /D1:1/U1:1/D1:1/U1:1/D1:1	DUDUDU DUD	10	105	85	52	6	2
A-PT	C/S0C/S0C/S0C/S0C/S0C/S0C	S	7	2	3	5	8	48
AD-PT	D:1/D0C:1/D0B-1:1	D	3	5	6	11	17	22
R-PT	0/S0/S0/S0/S0/S0/D1/S0/S0	SDS	9	2	8	11	12	130
RD-PT	0:1/S0:1/S0:1/S0:1/U1:1/S0:1/ S0:1/S0:1	SUS	8	2	8	9	55	103

Table 6.7: PTs that consistently progress upwards (highlighted in grey) or downwards through the ranking positions with the highest ranking position in the relevant first or last period

The two tables also show that there is *either* a large difference in the ranking positions of the PTs between period 1 and period 5 without the PTs making it into the highest ranking positions in the appropriate first or last period, *or* that the PTs start (or end) in the highest ranking positions in the appropriate first or last period but do not move significantly down (or up) the ranking positions. For example, the upward progressing A-PT in Table 6.6 moves from ranking position 629 in period 1 to ranking position 92 (out of 511 rankings) in period 5 and therefore does not make it into the top 10% of ranking positions in period 5; whereas the upward progressing A-PT in Table 6.7 progresses to the highest ranking position in period 5 from ranking position 82 (out of 747 rankings) in period 1.

Because PTs in Table 6.6 compared to those in Table 6.7 show a greater degree of movement through the ranking positions across the periods, and show a greater range of intervallic movement

[12] An explanation of the PT encoding is provided in Section 4.4.4.

between the individual notes, they will be investigated further by looking at their musical context. Example 6.1 shows, in standard music notation, all eight PTs from Table 6.6 with the R-PT and RD-PT types starting on c^2, with all the durations starting with a crotchet.

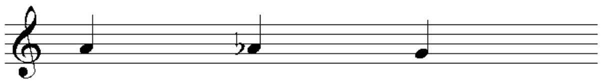

i) A-PT - A/D0A-1/D0G

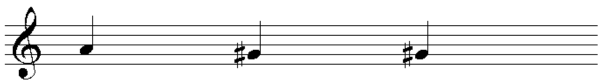

ii) AD-PT - A:1/D0G+1:1/S0G+1:1

iii) R-PT - 0/U1/U3/D2

iv) RD-PT - 0:1/D4:1/U3:1

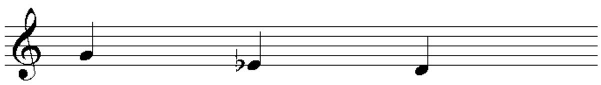

v) A-PT - G/D0E-1/D0D

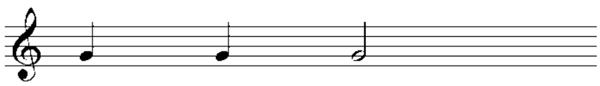

vi) AD-PT - G:1/S0G:1/S0G:2

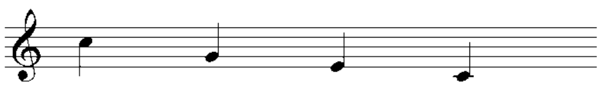

vii) R-PT - 0/D5/D3/D4

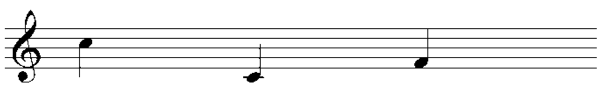

viii) RD-PT - 0:1/D12:1/U5:1

Example 6.1: Examples of the PTs in Table 6.6

Before looking at the location of the patterns in the music, it is important to determine if the PTs in Example 6.1 have a wide spread of instances across periods and composers and are not appearing in Table 6.6 simply because they are extremely frequent in a few compositions. Table 6.8 shows how many PTIs are in the dataset, how many movements and compositions each PT appears in, and how many composers have used each PT from Table 6.6.

Example Number	PTD	PT	Number			
			Of Total PTIs	In Movements	In Compositions	Of Composers
3.7 i	A-PT	A/D0A-1/D0G	349	130	66	18
3.7 ii	AD-PT	A:1/D0G+1:1/S0G+ 1:1	332	103	63	18
3.7 iii	R-PT	0/U1/U3/D2	251	97	62	18
3.7 iv	RD-PT	0:1/D4:1/U3:1	300	96	63	18
3.7 v	A-PT	G/D0E-1/D0D	447	143	73	17
3.7 vi	AD-PT	G:1/S0G:1/S0G:2	366	125	74	14
3.7 vii	R-PT	0/D5/D3/D4	656	182	92	14
3.7 viii	RD-PT	0:1/D12:1/U5:1	460	173	87	13

Table 6.8: Number of PTIs of each PT from Table 6.6

Table 6.8 shows that each PT has a large spread of PTIs across both compositions and composers (there are 111 pieces in total by 19 different composers). Therefore, the PTs occur in the majority of compositions by the majority of composers. This is interesting, because a lower figure for the number of composers would be expected. At the start (or end, depending on progression direction) of the PT's progression, it is relatively infrequent and should therefore appear in fewer compositions, and consequently composers' works, in appropriate the period. For upward progressing PTs, each appears in all but one composer's work, indicating widespread use in all periods. However, it should be remembered that there are only two composers in the first period, Haydn and Mozart, increasing the likelihood of the PTs appearing in all of the composers in the first period, despite the small number of PTIs.

So far it has been argued that the PTs in Table 6.6 may exhibit selection and replication on the grounds that they are either consistently moving upwards or downwards through the ranking positions across the periods, over a wide range of ranking positions, and over a wide range of compositions and composers. This now raises the question of where and how these PTs appear in real compositions. Do they appear in expositions, in codas, in trio sections, as part of cadences, etc.; in any of the instrumental parts; and are they a part of melodic interest, part of an accompaniment figure, part of a bass-line, part of a sequence, part of a parallel passage between instruments, etc. The location within the composition could show if the scenarios of PTIs have a bearing on whether the PT is selected and replicated. For example, it may be that the PTs with a high ranking position appear in similar scenarios, whereas those PTs at the bottom of the ranking positions appear in a wider variety of scenarios. If this situation is the case, then there would be supporting evidence for the scenarios of the PTIs having a bearing on the ranking positions, and by extension, on selection and replication.

Some of the PTs in Example 6.1 are, in their own way, quite interesting in terms of this study. For example, Example 6.1 i consists of three notes descending by semitones, which is a sequence of chromatic notes. Consequently, it would be expected that Example 6.1 i would be used less frequently in the earlier periods, when the traditional major-minor tonalities were central, and more frequently in the later periods, when composers started to break away from this system. This is exactly what the dataset is showing for Example 6.1 i, in that the PT progresses upwards through the ranking positions across the periods. A similar point can also be made for Example 6.1 iii, where the pitches cannot be mapped to any of the major scales. However, the other two PTs that progress upwards through the ranking positions, Example 6.1 ii and Example 6.1 iv, can both be mapped onto the traditional major-minor tonality.

Conversely, the PTs in Example 6.1 that progress downwards through the ranking positions across the periods (Example 6.1 v to Example 6.1 viii) can all be mapped onto a traditional major-minor tonality. For example, the PT in Example 6.1 vii is a descending arpeggio in a major key in root position. Again, the dataset mirrors what would be expected of such a PT, in that it is high in the ranking positions in the earlier periods, where the traditional major-minor tonalities were central to composers, and descends through the ranking positions to the later periods. The only possible exception to this phenomenon is the PT in Example 6.1 vi, which consists of three consecutive Gs. The majority of the examples in Example 6.1 conform to what would be expected in terms of the PTs' progressions through the ranking positions. However, there are exceptions in that some of the PTs are not easily identified as unambiguously belonging to the traditional major- minor tonalities, or as definitely *not* belonging to them. For example, Example 6.1 ii could belong to a number of different major or minor keys (such as A major or F# minor), but there is not enough information within the PT (given that there are only two different pitches) to say that it can definitely be mapped on to a specific major or minor tonality. Therefore, it is impossible to say if such examples accord with a specific directional movement within the ranking positions.

The discussion on the tonality of the PTs in Example 6.1 above ignores some of the complexities involved. For example, many chromatic passages can be integrated into a predominantly diatonic structure. However, some of the examples do show that some chromatic PTs are more dominant in the later than in the earlier periods, as would be expected.

Even though Example 6.1 i does not, on its own, conform to the traditional major- minor tonality, this does not mean to say that it will not have been used in compositions based firmly on that system. This is reflected in the dataset, which shows that Example 6.1 i is used by Haydn and Mozart. The first PTI of Example 6.1 i occurs in the first movement of Haydn's String Quartet op. 1 no. 1, where it is used as part of a downward chromatic accompaniment figure (see Example 6.2, where the PTI is highlighted in blue).

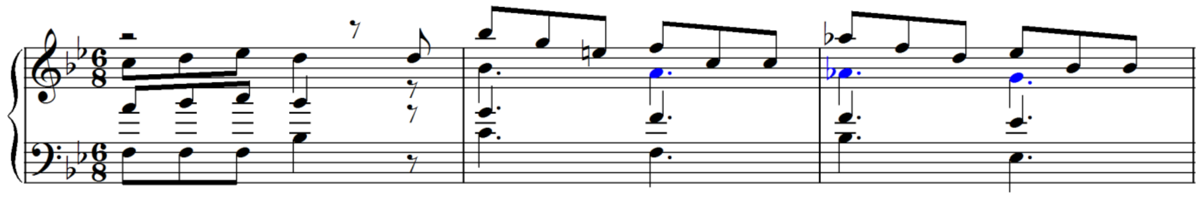

Example 6.2: Haydn String Quartet op. 1 no. 1 1st movement bb. 48-50

Example 6.2 shows that Example 6.1 i is part of a local chromatic passage in the key of B♭ major that can be interpreted as a repeated one-bar sequence consisting of two V^7 - I progressions in the keys of F major and E♭ major respectively. The PT occurs as part of the transitional passage between the two subject groups in the recapitulation section of the movement (the corresponding section in the exposition is worked differently, creating a different progression). However, the PT spans the two chord sequences. Therefore the first PTI of Example 6.1 i in the dataset could be considered a result of the employment of a sequence.

However, Example 6.1 i is not just restricted to the result of patterns formed within sequences, as in Example 6.2, but can also be a part of the main melodic material of a composition. Example 6.3 shows Example 6.1 i appearing as part of the repetition of the fugue subject in the fourth movement of Haydn's String Quartet op. 20 no. 2.

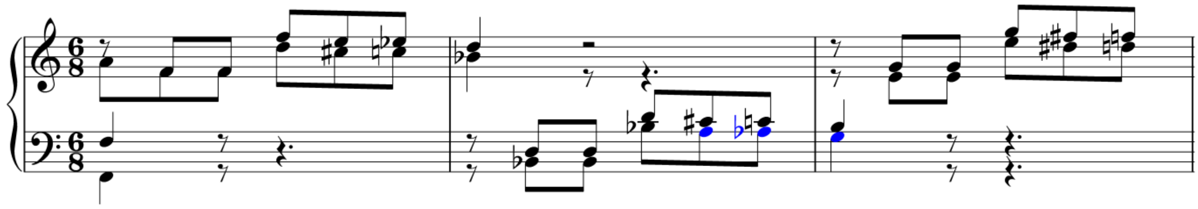

Example 6.3: Haydn String Quartet op. 20 no. 2 4th movement bb. 50-52

Example 6.3 shows Example 6.1 i occurring in the middle of a fugue movement, where the first six notes of the fugue subject are repeated in thirds alternating between the violins and the pairing of viola and cello (Example 6.4 gives the fugue subject). However, the PT cannot actually appear in the fugue subject in the home key of the movement (C major). Therefore, although this PT in its absolute pitch value form is not prominent in this movement (it only appears three times in this form), in its relative intervallic value form (i.e., three descending semitones) it pervades the movement owing to its prominence in the fugue subject.

Example 6.4: Fugue subject from Haydn's String Quartet op. 20 no. 2 4th movement

Mozart also gives the relative intervallic value of Example 6.1 i prominence in other compositions. Example 6.5 shows the pattern type in its absolute pitch value form in the second movement of Mozart's String Quartet K. 168.

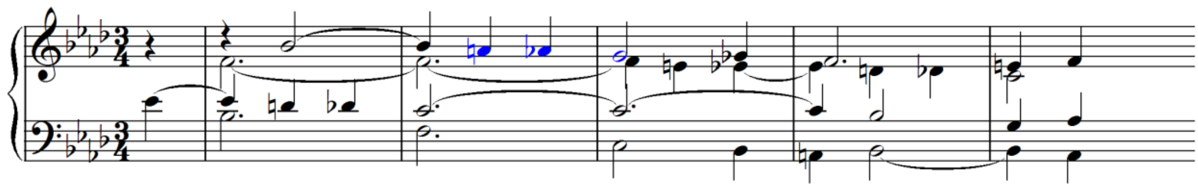

Example 6.5: Mozart String Quartet K. 168 2nd movement bb. 49-53

Example 6.5 shows part of the second subject from the recapitulation in the slow movement. The first subject is arguably in two halves (see Example 6.6), with the second half showing that it could form the basis of a transformation into the second subject as in the passage in Example 6.5, in that both passages begin with a downward stepwise progression. In the exposition the second subject is in a different key (the relative major) from the first subject, as would be expected, and therefore Example 6.1 i does not appear in its absolute pitch value form. However, it does appear in its relative intervallic value form in bars 15-16 in the first violin part. The PT also appears in a number of other places in its relative intervallic form. For example, in Example 6.5 the first three notes of the viola part follow the same relative intervallic form as Example 6.1 i due to the imitative nature of the movement.

Example 6.6: The first subject from Mozart's String Quartet K. 168 2nd movement

The examples above show that Example 6.1 i is found in different scenarios within compositions from period 1. The PT in its absolute pitch form is arguably not particularly prominent even though it appears in some main subjects, although its relative intervallic form has been shown as part of a fugue subject (Example 6.4 above).

An interesting point regarding the PTIs of Example 6.1 i shown above is that they are all preceded by a B♭. When looking at the dataset, 87% of PTIs of Example 6.1 i are preceded by a B♭, with the second most popular preceding pitch being B, albeit in only 7% of PTIs. The four-note PT consisting of Example 6.1 i preceded by a B♭ continuously progresses upwards through the ranking positions across the periods from position 503 in period 1 to position 128 in period 5. This progression mirrors that of Example 6.1 i, although Example 6.1 i has a greater difference in the ranking position between periods 1 and 5 (629 and 92 respectively), as well as achieving a higher position within the ranking positions in period 5. Therefore, there appears to be a link in terms of movement within the ranking positions between the PTs of Example 6.1 i and Example 6.1 i preceded by a B♭, with both showing a greater dominance in the ranking positions as they progress through the periods.

The PTIs of Example 6.1 i in period 2 are very similar to those in period 1. There are examples of the PT being used as part of an accompaniment figure, as part of melodic material, as well as in its relative intervallic form. Example 6.7 shows Example 6.1 i in the viola part of the first movement of Mozart's String Quartet K. 428.

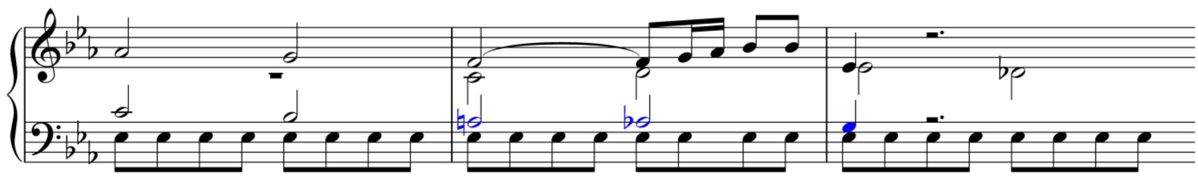

Example 6.7: Mozart String Quartet K. 428 1st movement bb. 153-155

This shows Example 6.1 i being used as part of a descending chromatic scale harmonic accompaniment figure during the coda section. This use Example 6.1 i also occurs in the codetta section at the end of the exposition in the second violin part (bb. 57-–59), but in a different key, i.e., in its relative intervallic values.

An example of Example 6.1 i being used as part of melodic interest occurs in the third movement of Haydn's String Quartet op. 74 no. 3, shown in Example 6.8. Here Example 6.1 i is used as part of the second phrase at the start of the second half of the Trio section. This PT is also used as part of the accompaniment to the main theme in the first half of the Trio section (see Example 6.9). Because the PT is appearing at the start of a second phrase and as part of the accompaniment to a main theme, it can be argued that this figure is an important aspect of this movement.

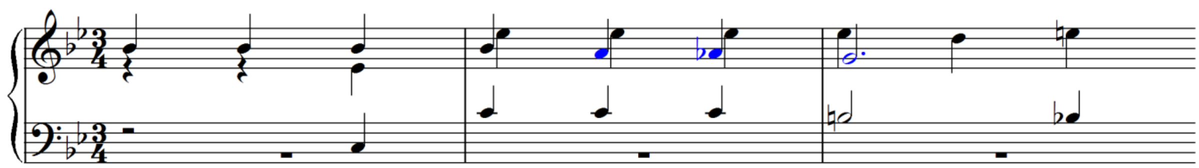

Example 6.8: Haydn String Quartet op. 73 no. 3 3rd movement bb. 47-49

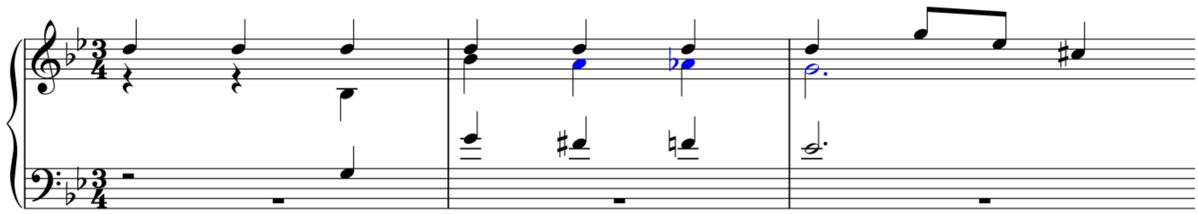

Example 6.9: Haydn String Quartet op. 73 no. 3 3rd movement bb. 35-37

Example 6.9 also shows that Example 6.1 i is being used in its relative intervallic form in the viola part in bar 36 as part of parallel-third part-writing with the second violin part using the absolute pitch form of Example 6.9. The PT in its relative intervallic form is also used in a number of other places in the Trio section, giving it arguably more prominence than other PTs in this section.

The examples from period 2 again show that Example 6.1 i is being used in different scenarios. Additionally, there is, as with the examples from period 1, the fact that the PT in its absolute pitch value form is not especially prominent, although its relative intervallic form does help to reinforce its prominence in some movements.

The remaining periods also show Example 6.1 i being used in different scenarios, with the PTIs being found as part of accompaniment figures, as part of melodic material, as a consequence of sequences, within development sections, within codas, etc., and with examples in all the different instrumental parts. For example, in period 3, Example 6.1 i is found in a descending bass line in the first movement of Beethoven's String Quartet op. 59 no. 1 (where it forms part of a descending two-bar sequence), and in the same period it appears in the introduction section of the opening movement of Beethoven's String Quartet op. 130 (see Example 6.10 and Example 6.11).

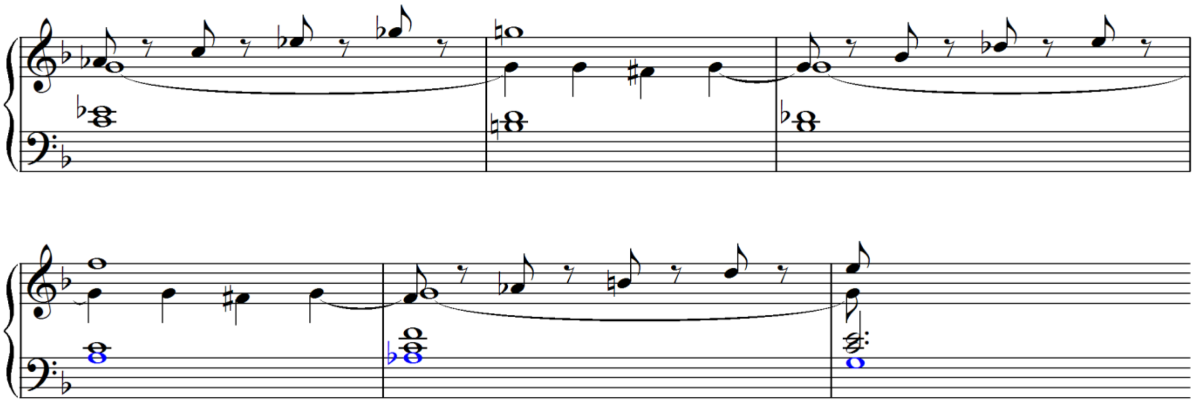

Example 6.10: Beethoven String Quartet op. 59 no. 1 1st movement bb. 53-58

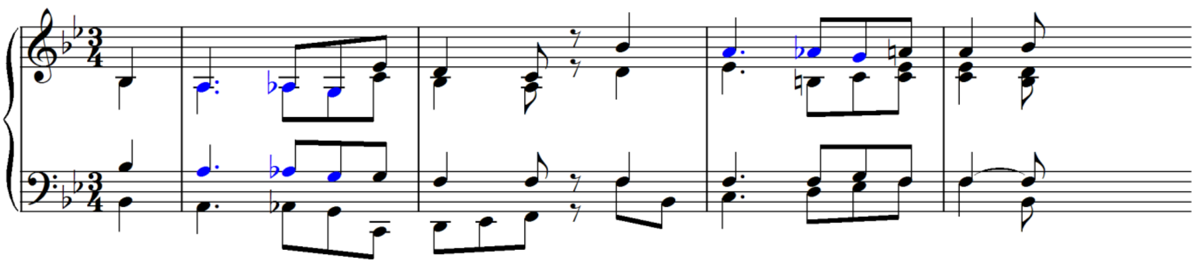

Example 6.11: Beethoven String Quartet op. 130 1st movement bb. 1-5

Example 6.1 i appears in a number of different scenarios across the periods. However, there are an increasing number of examples in the later periods where the PT is used as part of melodic material. For example, in period 5 the PT occurs as part of melodic material in works by Bartok and Shostakovich amongst others, with the PT also occurring in these compositions as part of accompaniment figures. Example 6.12 and Example 6.13 show the PT as part of the melodic material and as part of a chordal figure above a melodic line in the cello, respectively, in the final movement of Shostakovich's String Quartet no. 3. However, it should be noted that a relatively small number of samples were chosen, and consequently, may not reflect the dataset as a whole.

Example 6.12: Shostakovich String Quartet no. 3 5th movement bb. 77-80

Example 6.13: Shostakovich String Quartet no. 3 5th movement bb. 334-341

In the above examples, Example 6.1 i has been shown to occur as part of a number of different scenarios, i.e., as part of accompaniment figures, as part of melodic material, as a consequence of a sequence, as part of an opening slow section, as part of chordal passage, etc. It has also been shown that when the PT appears as part of the thematic material, its relative intervallic value equivalent PT can also be found. There is also some possible evidence for the PT being used more frequently as part of melodic material in the later periods, although this is based on a relatively small sample of examples.

The three other PTs in Example 6.1 that show an upwards progression through the ranking positions across the periods (i.e., Example 6.1 ii to Example 6.1 iv) also show that the PTs appear in a variety of scenarios across all the periods. The next set of examples show a sample of PTIs of the PTs in Example 6.1 ii to Example 6.1 iv, beginning with Example 6.14, which shows PTIs of Example 6.1 ii across all the periods.

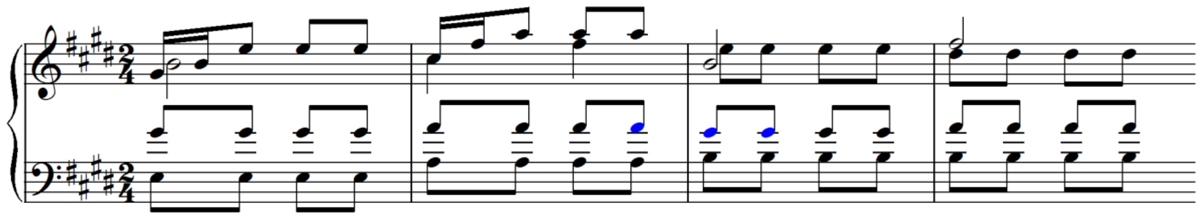

i) Haydn String Quartet op. 17 no. 1 4th movement bb. 198-201

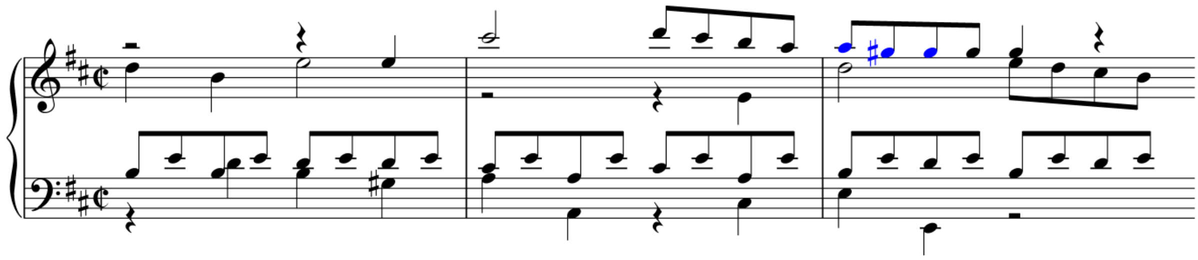

ii) Mozart String Quartet K. 575 4th movement bb. 33-35

iii) Beethoven String Quartet op. 59 no. 1 4th movement bb. 125-128

iv) Borodin String Quartet no. 1 1st movement bb. 432-438

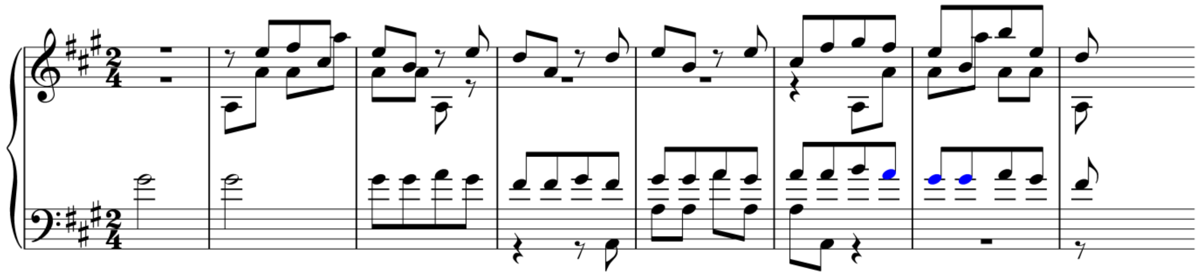

v) Sibelius String Quartet op. 56 2nd movement bb. 1-8

Example 6.14: Sample PTIs of Example 6.1 ii

The PTIs in Example 6.14 show Example 6.1 ii appearing in a number of different scenarios. Example 6.14 i, from Haydn's String Quartet op. 17 no. 1, shows Example 6.1 ii being used in period 1 as part of a repeated note accompaniment figure in the viola. In Mozart's String Quartet K. 575 from period

2, Example 6.1 ii is found in the fourth movement as shown in Example 6.14 ii as part of melodic material where the shape of the overall phrase is similar to part of the main theme of the movement (bb. 1-8), in that they both have a long note followed by a descending scale with the phrase ending with a falling semitone (bb. 3-4). Example 6.1 ii is being used here to vary the first four bars of the main theme, with the melodic material being split between violin 1 and violin 2.

Example 6.14 iii and Example 6.14 iv (from periods 3 and 4 respectively) show Example 6.1 ii as part of a repeated pattern of notes. Firstly, Example 6.14 iii, from Beethoven's String Quartet op. 59 no. 1, fourth movement, has Example 6.1 ii being used as part of a repeated chord section. Secondly, Example 6.14 iv, from Borodin's String Quartet no. 1, first movement, has Example 6.1 ii as part of a contrasting figure that alternates a continuous parallel octave semiquaver movement in the viola and cello parts with chords in the violin parts over longer notes in the viola and cello parts.

Finally, Example 6.14 v shows Example 6.1 ii being used as part of the main theme at the opening of the second movement of Sibelius' String Quartet op. 56 from period 5.

The next pair of examples (Example 6.15 and Example 6.16) show PTIs of Example 6.1 iii and Example 6.1 iv across all five periods.

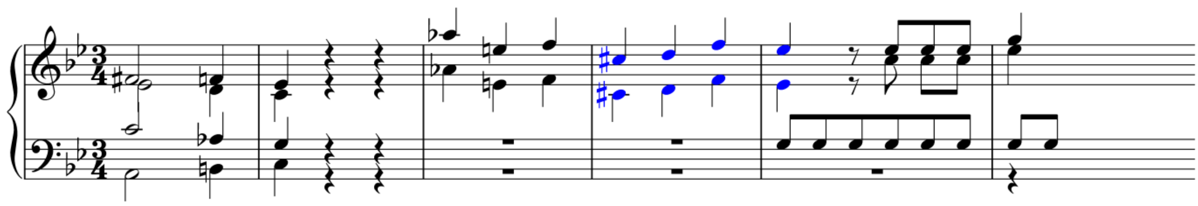

i) Mozart String Quartet K. 159 2nd movement bb. 136-141

ii) Haydn String Quartet op. 54 no. 2 4th movement bb. 56-60

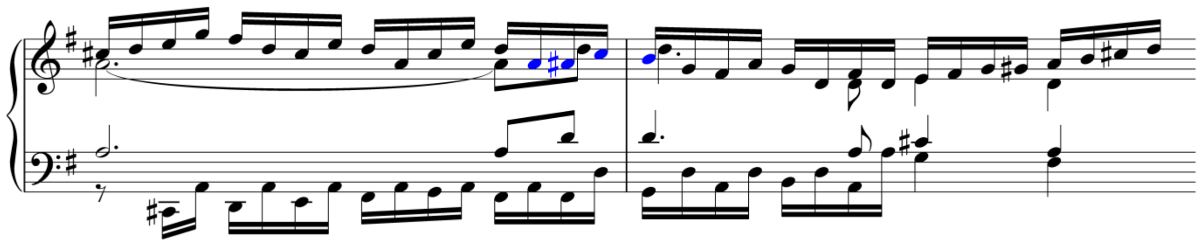

iii) Mendelssohn String Quartet op. 44 no. 2 3rd movement bb. 31-32

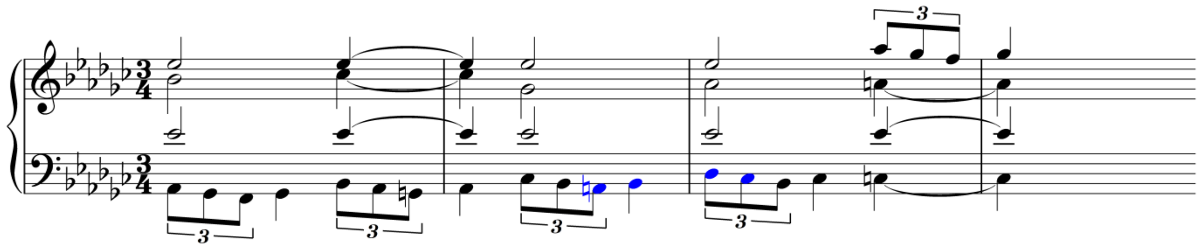

iv) Tchaikovsky String Quartet no. 3 1st movement bb. 562-565

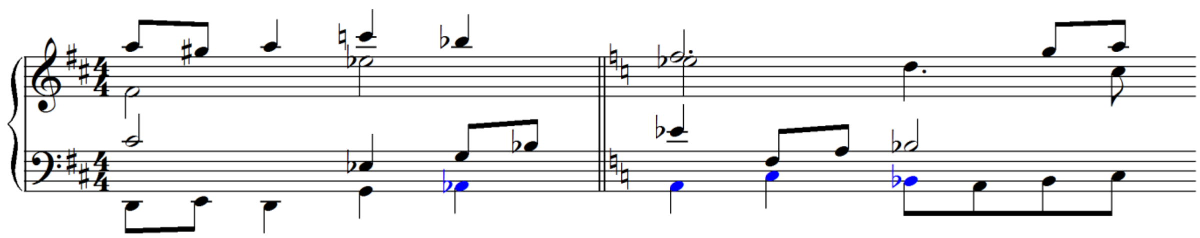

v) Prokofiev String Quartet no. 1 1st movement bb. 46-47

Example 6.15: Sample PTIs of Example 6.1 iii

i) Mozart String Quartet K. 173 3rd movement bb. 1-4

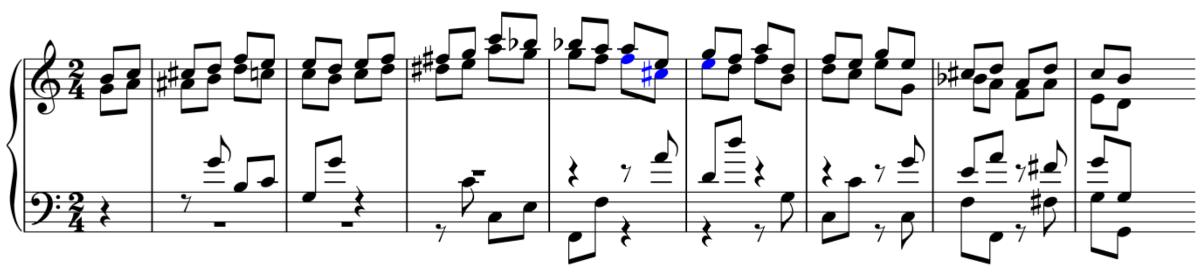

ii) Haydn String Quartet op. 54 no. 2 4th movement bb. 64-72

iii) Schumann String Quartet no. 1 2nd movement bb. 80-84

iv) Debussy String Quartet op. 10 2nd movement bb. 1-4

v) Janáček String Quartet no. 1 2nd movement bb. 1-4

Example 6.16: Sample PTIs of Example 6.1 iv

The examples of PTIs of Example 6.1 iii and Example 6.1 iv in Example 6.15 and Example 6.16, like those in Example 6.14, show a wide range of scenarios for the PTs. They range from being important aspects of the melodic material, such as Example 6.16 iv, where Example 6.1 iv forms part of the main theme of the second movement (and in this string quartet, the fourth movement as well), to part of an accompaniment figure, as in Example 6.16 i, where Example 6.1 iv forms part of the viola part.

There are other examples of the use of Example 6.1 iii and Example 6.1 iv as part of melodic material, such as Example 6.15 i, where Example 6.1 iii is used by the violins in as part of a parallel octave passage to lead into the next section of the movement; and Example 6.16 v, where Example 6.1 iv is used as part of the main opening melodic material in the viola part. Other examples where Example 6.1 iii and Example 6.1 iv are used as part of the melodic material include Example 6.16 ii, where Example 6.1 iv occurs due to the second violin following the first violin in thirds, and Example 6.15 ii, where Example 6.1 iii occurs as part of the first violin part, with the second violin part again following in thirds. An example of a bass-line figure occurs in Example 6.15 v where the cello part uses Example 6.1 iii as the movement changes key signature.

The placement across a change in the key signature in Example 6.15 v highlights an issue with the methodology, in that all possible patterns are generated, causing some PTIs to occur across boundaries such as structural breaks, time signature changes, key changes, etc. Because this example straddles two different key signatures, it also raises the question (together with PTs straddling changes of key not recorded by a key signature) of the importance of scale degrees within patterns, which is not addressed within the methodology.

Other PTIs of Example 6.1 iii are found in Example 6.15 iii, where the PT occurs as part of a continuous semiquaver line in the first violin part, and in Example 6.15 iv, where it occurs as a consequence of a sequenced triplet cello figure. The remaining PTI of Example 6.1 iv is in Example 6.16 iii, where the PT occurs as part of repeated melodic material.

So far, all the examples used in this section are of PTs that are continuously *ascending* through the ranking positions across the periods from Table 6.6. The next set of examples (Example 6.17 to Example 6.20) show those PTs that are continuously *descending* through the ranking positions across the periods from Table 6.6.

i) Haydn String Quartet op. 17 no. 3 3rd movement bb. 6-13

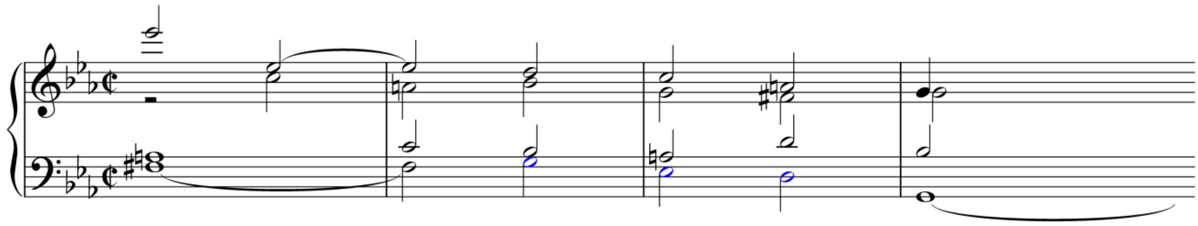

ii) Schubert String Quartet no. 1 1st movement bb. 6-9

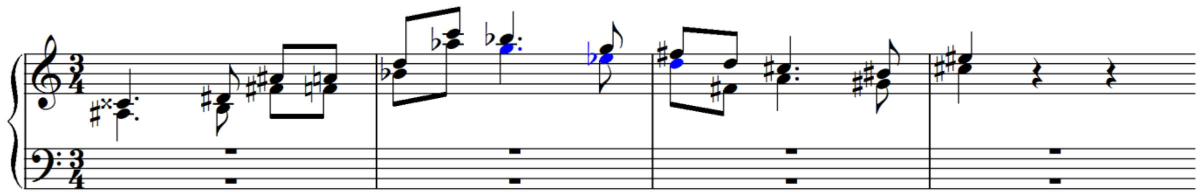

iii) Bartok String Quartet no. 1 2nd movement bb. 11-14

Example 6.17: Sample PTIs of Example 6.1 v in periods 1, 3 and 5 respectively

i) Haydn String Quartet op. 33 no. 3 4th movement bb. 9-15

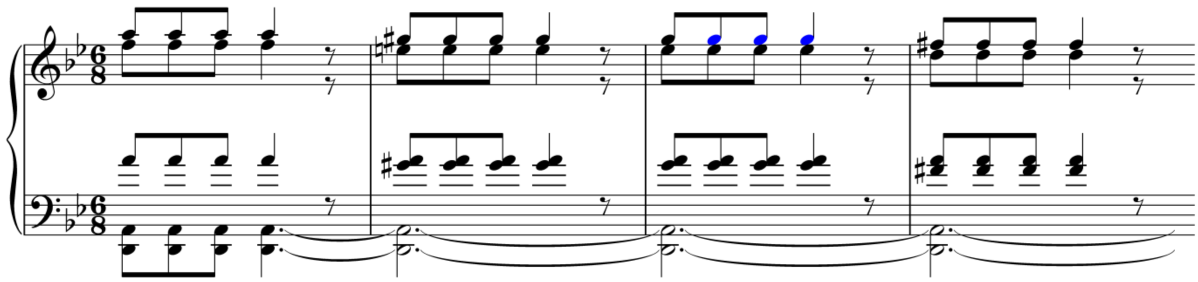

ii) Grieg String Quartet op. 27 4th movement bb. 193-196

Example 6.18: Sample PTIs of Example 6.1 vi in periods 2 and 4 respectively

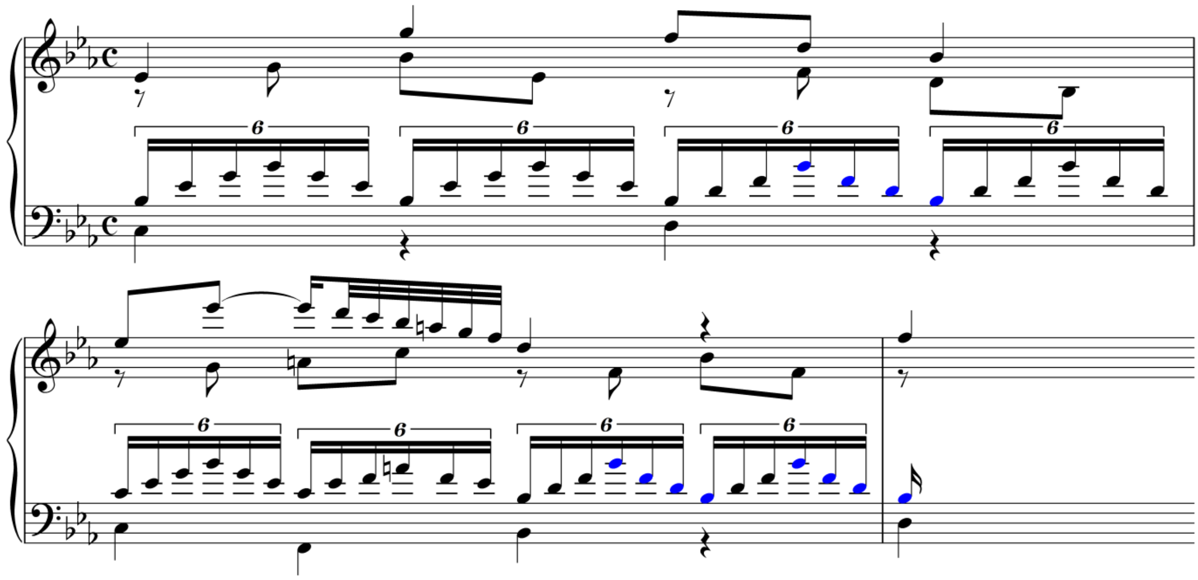

i) Haydn String Quartet op. 22 no. 2 2nd movement bb. 38-40

ii) Schumann String Quartet no. 3 3rd movement bb. 1-3

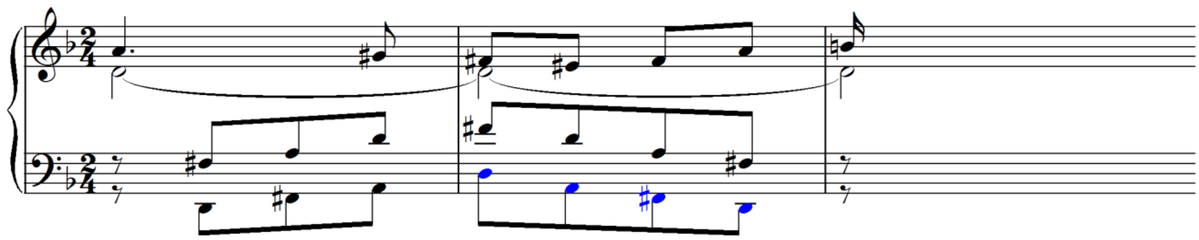

iii) Sibelius String Quartet op. 56 5th movement bb. 51-53

Example 6.19: Sample PTIs of Example 6.1 vii in periods 1, 3 and 5 respectively

i) Beethoven String Quartet op. 18 no. 1 4th movement bb. 5-8

ii) Tchaikovsky String Quartet no. 1 4th movement bb. 354-359

Example 6.20: Sample PTIs of Example 6.1 viii in periods 2 and 4 respectively

Example 6.17 to Example 6.20 show that, like the PTs in Example 6.1 that consistently progress upwards through the ranking positions across the periods, there are a wide variety of scenarios for the PTIs of PTs from Example 6.1 that consistently progress downwards through the ranking positions across the periods. These include PTIs as part of melodic material, such as in Example 6.17 i, where Example 6.1 v is part of the melodic line in the first violin part in the Minuet section of the movement, and in Example 6.19 ii, where Example 6.1 vii appears in the opening theme in the first violin part. There are also examples of the PTs as part of bass-line figures, such as in Example 6.20 i, where Example 6.1 viii is part of the cello line for the last three notes of the bass line during the cadence point at the end of a phrase.

Other PTIs of the PTs from Example 6.1 v to Example 6.1 viii appear in bass lines (as in Example 6.17 ii and Example 6.19 iii), as an arpeggiated accompaniment figure (as in Example 6.19 i), and as part of a repeated chord figure (as in Example 6.18 ii). They can also be shown in extensions to melodic fragments, as in Example 6.20 ii, where the PTI is an extension to an imitated phrase passed through the parts in the preceding bars. Example 6.18 i is a PTI of Example 6.1 vi in the first violin part as part of a rhythmic figure whilst the second violin and viola provide melodic movement in thirds. Finally, Example 6.17 iii shows a PTI of Example 6.1 v as part of a parallel-thirds passage between the first and second violin parts.

All the PTIs used in Example 6.2 to Example 6.20 have shown that the PTs in Example 6.1 occur in a wide variety of scenarios. However, all the PTs in Example 6.1 are either three- or four-notes long. The examples above show that the PTs in Example 6.1 are, on the whole, part of longer patterns or phrases. Therefore, do longer PTs also show a wide variety of scenarios in the compositions from the shorter PTs? The seven- and eleven-note PTs represent the average length and the longest length PTs respectively, so an investigation of these can be used to counterbalance the predominance of the shorter-length PTs in Table 6.6. Table 6.9 adapts the format of Table 6.6 for seven- and eleven -note PTs only, together with the corresponding example number in Example 6.21 shown below the table. Those rows highlighted in grey are for PTs that ascend through the ranking positions, whilst those that are not highlighted are for PTs descending through the ranking positions.

PTD	Example Number	PT	Shape Property	PTL	Ranking Position within Period				
					1	2	3	4	5
A-PT	3.27 i	E/U0F/D0E/U0F/D0E /U0F/D0E	UDUDUD	7	181	171	107	50	24
AD-PT		NONE		7					
R-PT	3.27 ii	0/U2/U2/U1/U2/U1/U2	U	7	247	215	206	173	12
RD-PT	3.27 iii	0:1/D1:1/D2:1/D1:1/U1:1 /U2:1/U1:1	DU	7	186	170	162	143	82
A-PT	3.27 iv	C/S0C/S0C/S0C/S0C/S0C /S0C	S	7	2	3	5	8	48
AD-PT	-	NONE	-	7	-	-	-	-	-
R-PT	3.27 v	0/S0/S0/S0/S0/S0/D7	SD	7	44	79	85	101	234
RD-PT	3.27 vi	0:1/S0:1/S0:1/D1:1/S0:1 /S0:1/S0:1	SDS	7	5	12	16	18	158
A-PT	-	NONE	-	11	-	-	-	-	-
AD-PT	-	NONE	-	11	-	-	-	-	-
R-PT	3.27 vii	0/S0/R0/S0/S0/R0/S0/S0 /R0/S0/S0	S	11	141	63	39	20	18
RD-PT	3.27 viii	0:1/U1:1/D1:1/U1:1/D1:1 /U1:1/D1:1/U1:1/D1:1 /U1:1/D1:1	UDUDUDU DUD	11	92	79	46	6	3
A-PT	-	NONE	-	11	-	-	-	-	-
AD-PT	-	NONE	-	11	-	-	-	-	-
R-PT	3.27 ix	0/S0/S0/S0/S0/S0/S0/S0 /D2/S0/S0	SDS	11	12	20	26	39	92
RD-PT	3.27 x	0:1/S0:1/S0:1/S0:1/S0:1 /S0:1/D1:1/S0:1/S0:1 /S0:1/S0:1	SDS	11	2	7	11	12	62

Table 6.9: Seven- and eleven-note PTs that consistently progress upwards or downwards through the ranking positions with the largest difference in ranking positions between the first and last periods

i) A-PT - E/U0F/D0E/U0F/D0E/U0F/D0E

ii) R-PT - 0/U2/U2/U1/U2/U1/U2

iii) RD-PT - 0:1/D1:1/D2:1/D1:1/U1:1/U2:1/U1:1

iv) A-PT - C/S0C/S0C/S0C/S0C/S0C/S0C

v) R-PT - 0/S0/S0/S0/S0/S0/D7

vi) RD-PT - 0:1/S0:1/S0:1/D1:1/S0:1/S0:1/S0:1

vii) R-PT - 0/S0/R0/S0/S0/R0/S0/S0/R0/S0/S0

viii) RD-PT - 0:1/U1:1/D1:1/U1:1/D1:1/U1:1/D1:1/U1:1/D1:1/U1:1/D1:1

ix) R-PT - 0/S0/S0/S0/S0/S0/S0/S0/D2/S0/S0

x) RD-PT - 0:1/S0:1/S0:1/S0:1/S0:1/S0:1/D1:1/S0:1/S0:1/S0:1/S0:1

Example 6.21: Examples of the PTs in Table 6.9

Table 6.9 shows some differences from the shorter three- and four-note patterns of Table 6.6. Firstly, not all of the combinations of PTL/PTDs in Table 6.9 have a PT that consistently progresses upwards or downwards through the ranking positions across all the periods. For example, there are no eleven-note A-PTs that consistently progress upwards or downwards through the ranking positions across all the periods. Likewise, both the seven- and eleven-note AD-PT have no PTs in Table 6.9. This is a consequence of there being far fewer ranking positions available for the eleven-note as against the three-note PTs, with 75 versus 747 ranking positions available respectively in period 1 for A-PTs. Therefore, there are fewer available ranking positions for the longer- than the shorter-note PTs to progress through across the periods.

Secondly, there are not such wide differences in the ranking positions between periods 1 and 5 in Table 6.9 as there are in Table 6.6. For example, the seven-note A-PT in Table 6.9 moves from ranking position number 181 in period 1 to position 24 in period 5, whereas the corresponding PT in Table 6.6 moves from ranking position number 629 in period 1 to position 92 in period 5. Again, this is due to there being fewer possible ranking positions for longer than shorter PTLs.

Example 6.21 highlights further differences between the PTs in Table 6.6 and Table 6.9. For example, most of the PTs in Example 6.21 have very little intervallic movement and where there is movement, it is either by minor or major second (apart from Example 6.21 v, where the first ten notes are the same followed by a descending perfect fifth), whereas five out of the eight PTs in Example 6.1 (iii, iv, v, vii and viii) have at least one interval greater than a major second. Also, eight out of the ten PTs in Example 6.21 (all except ii and iii) have either a repeated note or a repeated two-note group, whereas only two out of the eight PTs in Example 6.1 (ii and vi) have a repeated note.

Another difference between the PTs in Example 6.21 and Example 6.1 relates to tonality. All the PTs in Example 6.21 can be mapped onto a traditional major-minor tonality, depending on the scale degree matching, whereas the same cannot be said for all the PTs in Example 6.1. Even the PTs in Example 6.21 that consistently ascend through the ranking positions can be mapped on to a traditional major-minor tonality, depending on the scale degree matching. Although a major-minor tonality mapping does not preclude a PT from being able to ascend through the ranking positions, it is rather unexpected because the composers in the later periods tended to move away from the traditional major-minor tonalities of the earlier periods. By comparison, only one out of the four PTs in Example 6.1 (Example 6.1 ii) that consistently ascends through the ranking positions across the periods can be mapped onto a major-minor tonality depending on the scale degree matching.

There are also PTs in Example 6.21 that are similar to each other in terms of the number of different notes and their intervallic movement, both within their own and across the two different PTL groups. For example, Example 6.21 iv is similar to Example 6.21 v (both of which are seven-notes long) in that their first six notes are the same, with only the final note being different. An example where there is a similarity across the different PTLs is shown in Example 6.21 vi and Example 6.21 x, where both PTs begin with a repeated note, move down a semitone, and then repeat the new note to the end of the pattern. In fact, Example 6.21 vi can be found nested within Example 6.21 x, as is shown in Example 6.22.

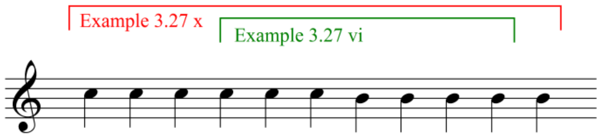

Example 6.22: Example 6.21 vi nested within Example 6.21 x

Finally, Example 6.21 vii is interesting in that it is the only PT in either Example 6.1 or Example 6.21 that includes rests. The PT itself can, with the help of accents in the correct places, be part of the classic 3/4 waltz accompaniment. This PT ascends through the ranking positions, meaning that there are more PTIs of it in relation to those of other PTs in the later periods than in the earlier periods, possibly reflecting its popularity in the late nineteenth-century.

The next set of examples (Example 6.23) shows PTIs of the seven-note PTs from Example 6.21.

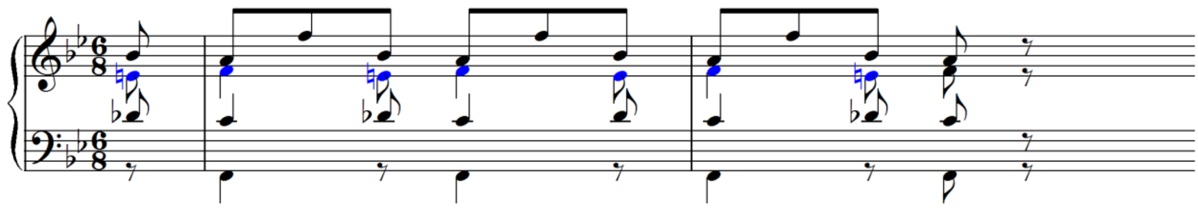

i) Haydn String Quartet op. 1 no. 1 1st movement bb. 38-40: a PTI of Example 6.21 i in period 1

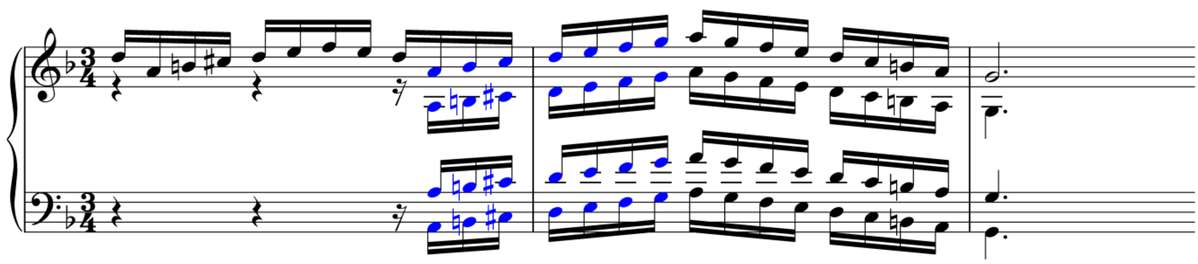

ii) Beethoven String Quartet op. 18 no. 1 1st movement bb. 80-82: a PTI of Example 6.21 ii in period 2

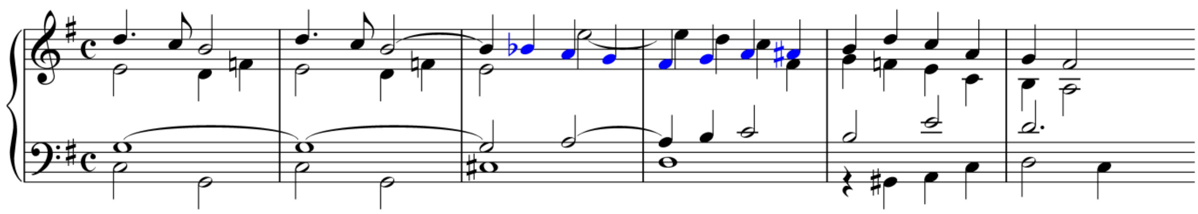

iii) Mendelssohn String Quartet no. 4 2nd movement bb. 57-62: a PTI of Example 6.21 iii in period 3

iv) Brahms String Quartet op. 51 no. 1 1st movement bb. 1-2: a PTI of Example 6.21 iv in period 4

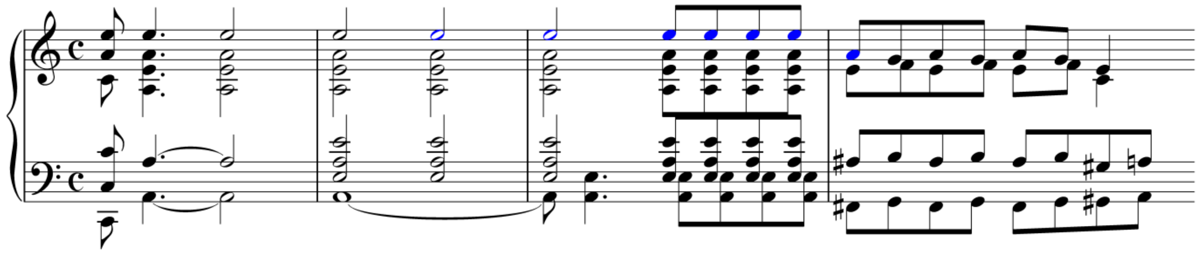

v) Bartok String Quartet no. 1 3rd movement bb. 304-307: a PTI of Example 6.21 v in period 5

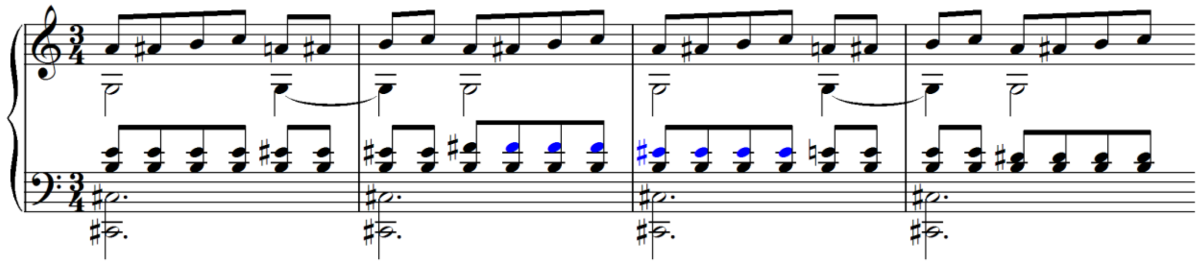

vi) Bartok String Quartet no. 2 2nd movement bb. 435-438: a PTI of Example 6.21 vi in period 5

Example 6.23: Sample PTIs of seven-note PTs from Example 6.21

Like the PTIs of the shorter PTLs from Example 6.1, the seven-note PTs from Example 6.21 occur in a wide variety of scenarios. For example, there is a PTI that is part of the melodic interest, such as Example 6.23 iii (where, interestingly, the PTI uses an enharmonic equivalent first and last note in the PT), and a PTI that is the top part of a chordal passage, as in Example 6.23 v. Another scenario is found in Example 6.23 ii, where Example 6.21 ii is part of a unison octave scale figure at the end of a transitional passage. The PTs from Example 6.21 can also be found as accompaniment figures, especially those with a number of repeated notes. For example, both Example 6.23 iv and Example 6.23 vi involve the PTIs as part of a repeated-note pattern. Finally, Example 6.21 i, appears as part of a repeated three-note figure in Example 6.23 i.

Unlike the shorter PTLs in Example 6.1, which appear in a variety of scenarios, the longer seven-note PTs in Example 6.21 tend, on the whole, to appear more in accompaniment figures rather than as part of the melodic interest. This change of scenarios for longer PTLs becomes even more prominent when looking at the eleven-note PTs from Example 6.21, as shown in Example 6.24.

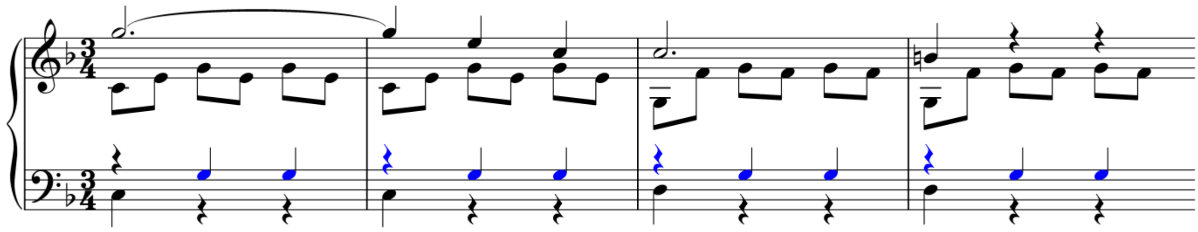

i) Mozart String Quartet K 158 1st movement bb. 12-15: a PTI of Example 6.21 vii in period 1

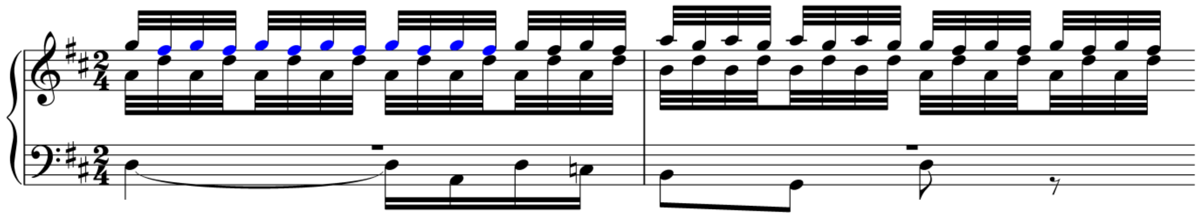

ii) Beethoven String Quartet op. 18 no. 5 3rd movement bb. 60-61: a PTI of Example 6.21 viii in period 2

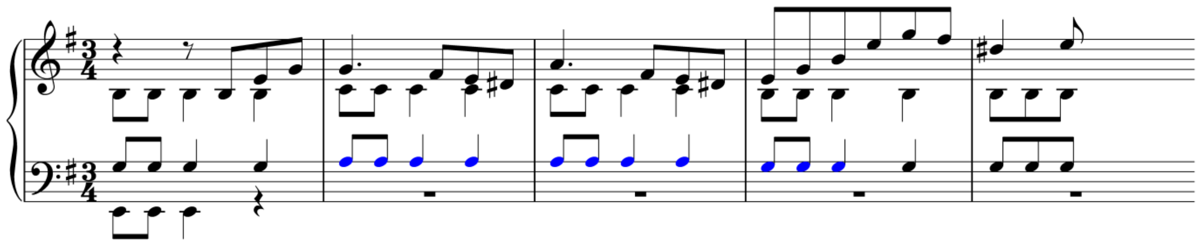

iii) Mendelssohn String Quartet no. 4 4th movement bb. 1-5: a PTI of Example 6.21 ix in period 3

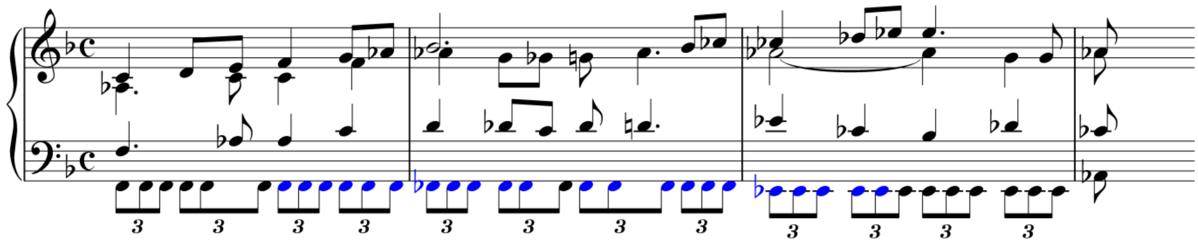

iv) Debussy String Quartet op. 10 1st movement bb. 49-52: a PTI of Example 6.21 x in period 4

Example 6.24: Sample PTIs of eleven-note PTs from Example 31

All the PTIs of eleven-note PTs from Example 6.21 in Example 6.24 involve the PTIs as part of accompaniment figures. These PTIs range from being part of a waltz-type accompaniment figure, as in Example 6.24 i, to being part of a repeated-note accompaniment figure, as in both Example 6.24 iii and Example 6.24 iv. Finally, Example 6.24 ii shows a PTI of Example 6.21 viii as part of a repeated two-note accompaniment figure over a cello melody.

6.2.3 Selection and Replication Summary

Some evidence has been shown for selection and replication by looking at the data *en masse*. The data has shown that there are certain PTs that are present in all five periods, although these are just a fraction of the total number of PTs. Of those PTs that are present in all five periods, there are a small number that progress either continuously upwards or continuously downwards through the ranking positions across the periods. Additionally, there are a number of PTs that do not appear in period 1 but then continuously progress upwards through the ranking positions across the periods. Likewise, there are PTs that continuously progress downwards through the ranking positions across the periods but that do not appear in period 5.

Because there are PTs that are present in the ranking positions across all the periods, as well as PTs that are appearing and disappearing from the ranking positions across the periods selection and replication may be taking place. If there were no selection and replication of the PTs, then the logical outcome would be that the vast majority of the PTs would appear in all the periods and would stay in roughly the same ranking positions. Also, there are a vast number of possible combinations of notes, and therefore of PTs, of which only a small proportion are represented within the dataset (although not all combinations would be musically viable). Again, this points to selection and replication taking place, in that only some of the PTs within the dataset (which represent only some of the possible combinations of notes), are appearing in all the periods. Additionally, the data clearly shows that the majority of PTs only appear in one of the five periods, pointing to these PTs not being successfully selected and replicated.

There is, however, a significant variation in the number of PTs appearing in just one out of the five periods across the different PTLs. For the shorter PTLs, the percentage of PTs that appear in just one period is far lower than is the case with the longer PTLs. This would suggest that there is less replication of the longer than the shorter PTs taking place, because there are far fewer longer PTs appearing in all five periods compared to those that appear in just one period.

For all PTLs, the number of PTs that appear in all five periods is significantly smaller than those PTs that appear in just one period. Additionally, of those PTs that continuously progress either upwards or downwards through the ranking positions across all the periods, there is an even smaller number that show a progression across at least 100 ranking positions between periods 1 and 5, i.e., that show a marked increase or decrease in ranking positions across the periods.

When looking at the spread across composers and their compositions of the PTs that either progress upwards or downwards through the greatest number of ranking positions across the periods, the data shows that these PTs occur in a range of string quartets and across a number of composers (although the dataset is not a complete record of every string quartet by every composer, it does cover 111 compositions from 19 different composers ranging from the classical period to the twentieth-century). It could be argued that those PTs that occur across a range of string quartets are just part of basic compositional techniques and are genre- or instrument-specific constraints. But those PTs that have the highest difference in the ranking positions between the first and the last periods are showing that they are either being used more, or used less across the periods and, by extension, are either being selected, or rejected, by composers.

The PTs that have the greatest difference between the first and last period ranking positions are all either three- or four-notes long. If these are being replicated and selected, then having just three or four notes in each PT does not accord with Miller's hypothesis of chunking information into groups of seven plus or minus two events (1956, p. 81). However, it does accord with his idea that we can easily remember four notes but then it becomes increasingly more difficult to remember five notes and upwards (1956, p. 81). There is the additional problem that there are far fewer longer than shorter PTs, meaning that the longer PTs do not have as much scope to progress through the ranking positions as the shorter ones.

The sampled PTIs from Example 6.1 have been shown in various scenarios within compositions. These range from melodic material from main themes to parts of accompaniment figures. Despite this range of scenarios, all the PTs were part of a longer material. This does not mean that there are no examples of a PT matching a whole phrase, rather that the examples taken did not indicate any. The placement of the PTs within compositions also varied. Some would be present at structural points, such as cadences at the end of sections, or within the main thematic sections; whilst others would appear in transitional sections, as part of a longer phrase, or even as a mixture of these scenarios.

Additionally, the PTs from Example 6.1 are varied in terms of mapping the pitches onto major-minor scale degrees. Some can definitely be mapped onto part of a traditional major-minor tonality, whilst others definitely cannot. There are PTs in Example 6.1 that progressed upwards through the ranking positions across the periods that cannot be mapped onto the traditional major-minor tonalities. Likewise, there were PTs in Example 6.1 that progressed downwards through the ranking positions across the periods that can be mapped onto traditional major-minor tonalities. This is to be expected, in that it shows the traditional major- minor tonalities are used less frequently in the later than in the earlier periods.

The examples of the longer PTLs in Example 6.21 show a different picture from that of the shorter PTLs in Example 6.1. Whereas there was a variety of scenarios within the shorter-length PTIs, the longer-length PTIs show a more restricted use. Additionally, the longer-length PTs in Example 6.21 tend to have less movement across the ranking positions than the shorter-length PTs, as well as being used more often in accompaniment figures than as part of melodic interest.

The fact that the longer PTs exhibit different properties to the shorter ones could be due to the latter being more memorable than the former, especially when considering Miller's ideas on the number of events we can remember. The fact that the longer PTs in Example 6.21 tend to be less complex in terms of intervallic movement than the shorter ones in Example 6.1 could be the result of the limitations of memory in remembering long, complex strings of events, as discussed in Section 2.5.1 above

6.3 Variation

A composer will tend to utilise material that reflects his or her own style of composition and this style is largely determined by the composer's historical context, including the stylistic, geographical and instrumental constraints of the time. As part of this process, a composer either uses existing patterns that they or other composers have used, or they create patterns that have never been used before. These created patterns could either have absolutely no connection with, or be variants to, existing patterns.

For the evolutionary process of variation to be found, it needs to be shown that there are patterns in the dataset that are ancestrally related. In other words, the dataset needs to show that a mutation from one pattern to another has taken place. The mutation would need to show evidence that it has either been generally adopted over time (i.e., that it is being selected and replicated), or that over time the mutation has faded from use (i.e., that it is *not* being selected and replicated). Therefore, if patterns can be shown to be mutations of previous patterns, and if the mutations can be shown to have been selected and replicated over time, then they show evidence of the evolutionary processes of selection, replication and variation, and can potentially provide further evidence for the existence of memes in music.

Distinguishing between completely new and mutated patterns is where the difficulty lies. How can it be shown that a new pattern is indeed a mutation of an earlier pattern? Also, if two patterns show signs of similarity in terms of their pitch and/or rhythmic elements, is this enough to establish an ancestral relationship between them?

6.3.1 Exploring Variation within the Dataset

This section will explore the evidence for the evolution process of variation in three ways. Firstly, it examines the number of PTs for three of the PTLs in each period to see if there is an increase in the number of PTs used. Secondly, it investigates the different sets of PT properties calculated (as described in Section 3.3 above) in each period to see if the number increases. Thirdly, it explores whether there are PTs in the later periods that have not appeared in the earlier periods and investigates the resultant PTs for similarity with other PTs in the earlier periods.

If there are more PTs of each PTL in each subsequent period, then there are PTs in the later periods that have not appeared in the earlier periods. Those PTs that have not appeared in the earlier periods could be descendants of those that do appear in earlier periods. Table 6.10 shows the number of different PTs that appear in each of the five periods for PTLs of three, seven and eleven notes.

PTL	PTD	Period				
		1	2	3	4	5
3	A-PT	5,481	8,879	11,011	9,863	8,144
	AD-PT	20,803	43,924	45,572	35,878	19,317
	R-PT	870	1,259	1,478	1,177	1,048
	RD-PT	7,716	15,717	17,422	12,961	8,026
7	A-PT	79,483	209,923	168,230	113,098	48,341
	AD-PT	86,981	238,628	184,574	121,736	51,276
	R-PT	58,511	144,972	126,228	84,349	38,266
	RD-PT	72,539	195,669	155,608	99,840	43,878
11	A-PT	96,706	270,938	204,278	133,626	55,310
	AD-PT	98,150	277,022	208,848	136,820	56,687
	R-PT	86,208	239,824	184,778	117,036	50,266
	RD-PT	88,706	250,129	191,883	122,190	52,350

Table 6.10: Number of PTs for each PTD in all the periods

Table 6.10 does not show an increase in the number of different PTs across all periods for any of the PTL/PTD combinations. The number of three-note PTs peaks in period 3 whereas the number of seven- and eleven-note PTs peaks in period 2 for all the PTDs. This roughly parallels the number of pieces and number of notes within each period, which peaks in period 2 before falling continuously across the periods to period 5 (see Table 6.11). This matching of the peaks between Table 6.10 and Table 6.11 could be used to claim that the number of different PTs across the periods is not the result of variation of existing patterns but is simply a natural consequence of the number of pieces and notes in each period within the dataset.

Period	Number of Pieces	Number of Notes
1	25	148,819
2	44	418,444
3	21	331,822
4	13	224,882
5	8	98,054

Table 6.11: Number of pieces and notes within each period

Although Table 6.10 does not show an increase in the values for PTs in every subsequent period, it can still be argued that there is some evidence for variation. For the three-note PTL, there is an increase in the number of PTs between periods 2 and 3 across all the PTDs that does not correspond with the number of pieces or notes in each period from Table 6.11. Therefore, this increase cannot be explained away by the number of pieces and notes. One explanation could be that a greater diversity of PTs are being used by composers. This explanation then raises the question of where these extra PTs come from.

Looking at the overall number of PTs within each period does not show how many are repeated across more than one period; it only shows how many there are in each period. However, there could be a set of unique PTs used in each period. Table 6.12 builds on Table 6.10 by showing how many PTs existed in at least one, or were not used in any, of the previous periods.

Table 6.12 shows that the majority of seven- and eleven-note PTs in each of the periods have no PTIs in any of the relevant previous periods. This is especially true for the eleven-note PTs, where at least 95% of them in any of the period/PTD combinations have not been used in any of the previous periods. However, for the three-note PTs it is only in periods 2 and 3 where most of the PTs have not been used before. This shows a clear difference in the number of pre-existing PTs between shorter three-note and longer eleven-note PTs.

PTL	PTD	New or Reused	Period 2		Period 3		Period 4		Period 5	
3	A-PT	R	4,576	52%	6,931	63%	7,505	76%	6,151	76%
		N	4,303	48%	4,080	37%	2,358	24%	1,993	24%
	AD-PT	R	11,216	26%	17,800	39%	17,707	49%	10,067	52%
		N	32,708	74%	27,772	61%	18,171	51%	9,250	48%
	R-PT	R	792	63%	1,104	75%	1,062	90%	976	93%
		N	467	37%	374	25%	115	10%	72	7%
	RD-PT	R	4,995	32%	8,335	48%	8,071	62%	5,240	65%
		N	10,722	68%	9,087	52%	4,890	38%	2,786	35%
7	A-PT	R	12,727	6%	20,816	12%	14,566	13%	6,542	14%
		N	197,196	94%	147,414	88%	98,532	87%	41,799	86%
	AD-PT	R	5,073	2%	8,298	4%	5,562	5%	2,796	5%
		N	233,555	98%	176,276	96%	116,174	95%	48,480	95%
	R-PT	R	16,695	12%	27,711	22%	20,678	25%	9,785	26%
		N	128,277	88%	98,517	78%	63,671	75%	28,481	74%
	RD-PT	R	5,883	3%	10,081	6%	7,279	7%	3,535	8%
		N	189,786	97%	145,527	94%	92,561	93%	40,343	92%
11	A-PT	R	2,420	1%	4,452	2%	3,006	2%	1,333	2%
		N	268,518	99%	199,826	98%	130,620	98%	53,977	98%
	AD-PT	R	1,499	1%	2,534	1%	1,443	1%	621	1%
		N	275,523	99%	206,314	99%	135,377	99%	56,066	99%
	R-PT	R	3,942	2%	6,803	4%	4,820	4%	2,423	5%
		N	235,882	98%	177,975	96%	112,216	96%	47,843	95%
	RD-PT	R	2,041	1%	3,395	2%	2,238	2%	1,206	2%
		N	248,088	99%	188,488	98%	119,952	98%	51,144	98%

R = PTs reused from at least one previous period
N = PTs not used in any of the previous periods

Table 6.12: Number of PTs that have either been used or not used in any of the previous periods

The fact that the majority of PTs in each period have not been used before could provide evidence for variation. This on the basis that these new PTs in each period may have originated from existing PTs. However, the fact that for both the seven- and eleven-note PTs the difference between the numbers of new patterns compared to existing patterns in each of the periods is so large (especially for the eleven-note PTs) could be used to argue that there is no variation taking place. The argument would be that if variation does exist, then there should be more previously used than new PTs, because surely composers would, on the whole, use PTs already in existence.

In Section 5.4 it was shown that there are more possible longer than shorter-length PTs. This difference in the number of PTs could provide an explanation for both the dominance of new PTs (a large set of possible PTs for a composer to select from), and for the difference in the ratio between new and existing PTs occurring in the eleven- and three-note PTs (the result of a larger set of eleven-note than three-note possible PTs for a composer to select from).

Looking at the number of PTs within each period and whether they are new or existing PTs has shown that there is some possibility that variation has taken place. However, the figures do not indicate anything definite, because there is not a continual increase in the number of PTs across the periods, nor is there any concrete evidence that any of the new PTs are variations of existing PTs.

The next stage in addressing this problem is to examine whether PTs become more varied over the five periods, using the PT properties specified in Section 3.3 above for determining similarity. If there is an increase in the number of different values for each PT property, and in the range of each PT property value across the periods, then this would indicate that the PTs are becoming more varied. For a PT to be a variant of an existing PT, then the PT properties of the existing PT will possibly change, creating a new set of PT properties. Therefore, if there is an increase in the number of different sets of PT properties, then variation could have taken place. Additionally, if music has become more complex over time, it would be expected that the PT property values would become increasingly varied over time creating more complex sets of properties. Table 6.13 shows the number of different PT property values within each period.

	Pitch Properties			Duration Properties		
Period	Shape	Centre	High/Low	Shape	Centre	High/Low
1	1,679	400	42	1,601	9,947	1,596
2	2,154	480	46	1,834	23,250	3,330
3	2,301	466	48	1,837	21,804	2,937
4	2,214	450	46	1,939	17,721	2,668
5	1,595	445	43	1,552	19,272	2,865

Shape = The overall shape of the pitch or duration
Centre = The average pitch or duration of the notes
High/Low = The distance between the highest and lowest pitches, or between the longest and shortest duration
The calculations for the creation of the properties are explained in Section 4.6.7.

Table 6.13: Number of distinct values for each of the pattern type properties within the dataset

Table 6.13 shows a peak for the number of different pitch-pattern properties during either period 2 or 3 before continuously falling through the periods to period 5. This roughly parallels the figures for the number of PTs in each period, as shown in Table 6.10 (for three-note PTs the number of PTs peaks in period 3, whereas both the seven- and eleven-note PTs peak in period 2), and likewise the figures for the number of pieces and notes in each period, as shown in Table 6.11. However, the figures for the duration properties show a different picture, in that the progression of the values across the periods does not follow the same lines as for the pitch-pattern properties. It could be expected that if the music of the later periods is more complex (in terms of intervallic and rhythmic structures), then there should be a correspondingly greater complexity of both pitch *and* duration properties in the later periods. However, the figures suggest that (once the number of pieces and notes in each period has been taken into account) there is a similar level of complexity across the periods in terms of pitch, and a *greater* level of complexity in terms of duration in the later periods. Therefore, this would suggest that if there is a greater complexity in later music, it is coming more from the rhythmic than the pitch element of the music.

To investigate the difference between the pitch and rhythmic properties further, Table 6.14 and Table 6.15 show the number of new property values compared to the number of existing property values in each period for the pitch and rhythmic elements respectively.

PTL	Property Type	New or Reused	Period							
			2		3		4		5	
3	High/Low	N	4	10%	1	2%	-	0%	-	0%
		R	37	90%	39	98%	38	100%	38	100%
	Centre	N	11	15%	1	1%	2	3%	1	1%
		R	64	85%	72	99%	67	97%	69	99%
	Shape	N	-	0%	-	0%	-	0%	-	0%
		R	9	100%	9	100%	9	100%	9	100%
7	High/Low	N	5	11%	-	0%	1	2%	-	0%
		R	41	89%	45	100%	44	98%	43	100%
	Centre	N	16	19%	2	2%	2	2%	1	1%
		R	69	81%	80	98%	80	98%	80	99%
	Shape	N	-	0%	-	0%	-	0%	-	0%
		R	189	100%	189	100%	189	100%	189	100%
11	High/Low	N	4	9%	2	4%	1	2%	-	0%
		R	42	91%	46	96%	45	98%	43	100%
	Centre	N	12	14%	4	5%	-	0%	1	1%
		R	72	86%	81	95%	83	100%	81	99%
	Shape	N	655	31%	282	12%	141	6%	36	2%
		R	1,483	69%	2,006	88%	2,051	94%	1,419	98%

Table 6.14: Number of pitch property values that have either been used or not used in any of the previous periods

PTL	Property Type	New or Reused	Period							
			2		3		4		5	
3	High/Low	N	194	57%	113	33%	94	29%	108	31%
		R	144	43%	231	67%	229	71%	236	69%
	Centre	N	325	62%	211	39%	154	32%	157	33%
		R	199	38%	335	61%	322	68%	316	67%
	Shape	N	-	0%	-	0%	-	0%	-	0%
		R	9	100%	9	100%	9	100%	9	100%
7	High/Low	N	1,134	64%	578	37%	485	34%	592	37%
		R	643	36%	982	63%	962	66%	995	63%
	Centre	N	2,604	72%	1,647	49%	1,071	40%	1,377	48%
		R	994	28%	1,686	51%	1,581	60%	1,491	52%
	Shape	N	7	4%	1	1%	1	1%	-	0%
		R	177	96%	180	99%	184	99%	179	100%
11	High/Low	N	1,991	65%	1,062	40%	766	33%	899	38%
		R	1,059	35%	1,585	60%	1,587	67%	1,468	62%
	Centre	N	4,080	74%	2,561	50%	1,593	41%	1,958	47%
		R	1,424	26%	2,554	50%	2,327	59%	2,189	53%
	Shape	N	398	22%	178	10%	162	9%	34	2%
		R	1,394	78%	1,619	90%	1,709	91%	1,444	98%

Table 6.15: Number of duration property values that have either been used or not used in any of the previous periods

Both Table 6.14 and Table 6.15 show that the majority of property values in both the pitch and duration element are not new for each period, although the figures for new values in the duration element are on the whole greater than for the pitch element. However, there are more new than existing values in the High/Low and Centre duration properties in period 2 for all PTLs.

The fact that the pitch element property values are not increasing substantially across the periods reinforces the idea, derived from the data in Table 6.13, that the pitch element of PTs is not increasing greatly in complexity over time. Likewise, the fact that the duration element shows a greater use of new pattern property values compared to the pitch element suggests that it is the former that is increasing in complexity. However, the figures for duration show that there are fewer new values being created in the later periods than in period 2.

Investigating the PT property values has not revealed much evidence for variation. The data has shown that there is some difference between the periods when looking at the PT property values, but there is not the consistent increase in the values across the periods. What has been seen is that there is a difference between the pitch and the duration properties of PTs, with the latter arguably becoming more complex across the periods. This increased complexity does suggest the possibility of variation having taken place, in that the more complex PTs being created in later periods could be variants of simpler PTs from earlier periods.

The final investigation in this section is to examine the PTs that appear for the first time in period 3, following their paths across the ranking positions in the subsequent periods, and to determine if they have any possible ancestors in the first two periods. If it can be shown that there are possible ancestors to some of the new PTs appearing in period 3, and that these new PTs then show selection and replication taking place, it is possible that these variant PTs could be considered potential candidates to be memes in music.

In order to show the possibility of selection and replication of the PTs new in period 3, Table 6.16 shows the number of new PTs that first appeared in period 3, the number subsequently rising and falling through the ranking positions, and the remaining number that fit into none of the previous categories used in this table for three-, seven- and eleven-note PTs.

PTL	PTD	Nº of PTs New in Period 3	Only Appear in Period 3		Subsequent Ups		Subsequent Downs		Remainder	
3	A-PT	4,080	2,655	65.1%	387	9.5%	0	0.0%	1,038	25.4%
	AD-PT	27,772	23,931	86.2%	497	1.8%	0	0.0%	3,344	12.0%
	R-PT	374	227	60.7%	35	9.4%	0	0.0%	112	29.9%
	RD-PT	9,087	7,416	81.6%	259	2.9%	0	0.0%	1,412	15.5%
7	A-PT	147,414	143,634	97.4%	150	0.1%	0	0.0%	3,630	2.5%
	AD-PT	176,276	174,613	99.1%	71	0.0%	0	0.0%	1,592	0.9%
	R-PT	98,517	93,450	94.9%	332	0.3%	0	0.0%	4,735	4.8%
	RD-PT	145,527	143,444	98.6%	128	0.1%	0	0.0%	1,955	1.3%
11	A-PT	199,826	198,882	99.5%	37	0.0%	0	0.0%	907	0.5%
	AD-PT	206,314	205,820	99.8%	18	0.0%	0	0.0%	476	0.2%
	R-PT	177,975	176,588	99.2%	101	0.1%	0	0.0%	1,286	0.7%
	RD-PT	188,488	187,794	99.6%	70	0.0%	0	0.0%	624	0.3%

Table 6.16: Number of three-, seven- and eleven-note PTs first appearing in period 3, and their subsequent progression through the ranking positions

Table 6.16 (like Table 6.12) shows that the longer the PTL, the greater the number of new PTs. Of the PTs new in period 3, there are some that progress upwards through the ranking positions across periods 4 and 5. Once again, there is a difference in the figures between the three- and eleven-note PTs, with a greater number of the former progressing upwards through the ranking positions compared to the latter. However, none of the PTs new in period 3 shows a continuous progression downwards through the ranking positions across periods 4 and 5 for any of PTL shown in Table 6.16.

The 'remainder' figure shows the number of new PTs that do not progress continuously upwards or downwards through the ranking positions across periods 4 and 5.

The fact that there are no PTs new in period 3 that continuously progress downwards through the ranking positions across periods 4 and 5 is rather unexpected, because it should be possible for a variant PT to be created in the third period that is not selected and replicated as frequently in the subsequent periods. However, it could be the case that the majority of the PTs new in period 3 are not achieving a sufficiently high enough ranking position to allow for subsequent falls in their ranking positions across periods 4 and 5. Table 6.17 shows the number of PTs new in period 3 that have the lowest ranking position in that period (i.e., that appear at the bottom of the ranking tables for period 3).

PTL	PTD	Number of PTs New in Period 3	PTs New in Period 3 in the Bottom Ranking Position	
			Number	Percentage
3	A-PT	4,080	2,727	67%
	AD-PT	27,772	20,855	75%
	R-PT	374	229	61%
	RD-PT	9,087	6,376	70%
7	A-PT	147,414	122,272	83%
	AD-PT	176,276	148,508	84%
	R-PT	98,517	77,659	79%
	RD-PT	145,527	116,645	80%
11	A-PT	199,826	175,698	88%
	AD-PT	206,314	182,659	89%
	R-PT	177,975	149,662	84%
	RD-PT	188,488	159,845	85%

Table 6.17: Number of PTs new in period 3 having the lowest ranking position

Table 6.17 clearly shows that the vast majority of PTs new in period 3 have the lowest possible ranking position. This means that these new PTs have only one PTI in the whole of period 3 and are therefore unlikely to progress any further down the ranking positions in the subsequent periods.

Even the figures for the number of PTs progressing upwards through the ranking positions in Table 6.16 can be misleading when looking at the number of ranking positions for each period. This is because there are a different number of available ranking positions for each period within each PTL/PTD combination, as shown in Table 6.18.

PTL	PTD	Period				
		1	2	3	4	5
3	A-PT	747	1,318	1,222	1,003	511
	AD-PT	593	1,214	1,008	778	394
	R-PT	392	501	538	492	367
	RD-PT	615	1,086	954	769	439
7	A-PT	193	493	295	141	111
	AD-PT	138	378	198	102	78
	R-PT	372	780	632	396	234
	RD-PT	241	560	472	262	162
11	A-PT	75	163	91	70	44
	AD-PT	67	136	71	53	31
	R-PT	148	323	255	175	93
	RD-PT	123	261	210	118	62

Table 6.18: Number of available ranking positions for each period

Table 6.18 shows that there are different numbers of ranking positions across the PTLs, PTDs and periods. For example, the three-note A-PT shows that period 5 has fewer ranking positions than period 4, which in turn has fewer ranking positions than period 3 (this is in fact true for all PTL/PTD combinations). This will obviously have an effect when examining progression within the ranking positions across the subsequent periods for any PT in period 3.

There are, however, PTs that first appear in period 3 and the origin of these will now be explored. Table 6.19 shows the total number of PTs in period 2 together with the number of PTs new in period 3, and the number of PTs from period 2 that could be an antecedent of one of the PTs new in period 3. These antecedent figures are calculated by comparing the pattern properties of Shape, HighLow and PitchCentre (as explained in Section 4.6.7) of the PTs new in period 3 with those of all the PTs in period 2, using an exact criterion match between the pattern properties.

PTL	PTD	Total Nº of PTs in Period 2	Nº of PTs New in Period 3	Possible Antecedent PTs in Period 2	
3	A-PT	8,879	4,080	8,332	94%
	AD-PT	43,924	27,772	31,373	71%
	R-PT	1,259	374	209	17%
	RD-PT	15,717	9,087	2,020	13%
7	A-PT	209,923	147,414	200,397	95%
	AD-PT	238,628	176,276	81,189	34%
	R-PT	144,972	98,517	134,569	93%
	RD-PT	195,669	145,527	51,349	26%
11	A-PT	270,938	199,826	217,991	80%
	AD-PT	277,022	206,314	50,577	18%
	R-PT	239,824	177,975	189,765	79%
	RD-PT	250,129	188,488	37,904	15%

Table 6.19: Number of possible antecedents of new period 3 PTs in Period 2

Table 6.19 shows that, across all the PTL/PTD combinations, there are PTs new in period 3 that have exact matches in terms of pattern properties with some PTs from period 2. It should be noted that this table does not map the PTs in the two periods against each other; rather it simply shows the number of PTs that *can* be mapped. This means that there may be PTs new in period 3 that do not have any antecedents in period 2. It also means that there can be multiple matches between the two periods; i.e., a new PT in period 3 may have multiple possible antecedents in period 2, and a PT in period 2 may match a number of different PTs new in period 3 (a *many-to-many* relationship, in database terms). What the table does show is that there are, according to the similarity algorithm used, a number of possible matches between PTs across the two periods that could possibly indicate evidence of some form of variation relationship.

The percentages of the possible variants show a mixed picture across the PTLs and PTDs. For example, three-note PTs have the largest percentage of possible antecedents for A-PTs and AD-PTs (both absolute pitch value PTDs) whereas the seven- and eleven-note PTs have the largest percentage of possible antecedents for A-PTs and R-PTs (both PTDs without duration). The large percentage figures for the three-note A-PTs and AD-PTs could be an outcome of the methodology, in that a large number of the absolute pitch value PTs will have the same pattern properties. For example, a major ascending triad in root position can have a number of different pitch values (e.g., C-E-G and D-F#-A), creating a number of different PTs, but all of these PTs will have exactly the same pattern properties of 'U' for Shape, 7 for HighLow and 3.5 for PitchCentre. Three-note PTs also have fewer combinations for the duration of the notes, which could explain the high percentage value of the three-note AD-PTs.

If the PTs new in period 3 are descendants of PTs in period 2, then these new PTs need to also exhibit selection and replication in the subsequent periods in order to support their qualification as potential memes. Table 6.16 above showed that, of the PTs that first make an appearance in period 3 and then appear in both the subsequent periods, the vast majority move consistently upwards through the ranking positions, with none descending consistently through the ranking positions in the subsequent periods. As explained in Section 6.2.1 above, this movement through the ranking positions shows possible evidence for replication and selection. Once again, there is the problem of the number appearing in the later periods being smaller than the number of PTs appearing in the earlier periods, which can have an effect on the figures because there are fewer ranking positions available in the later periods. Consequently, the PTs that first appear in period 3 and then appear in subsequent periods are more likely to rise through the ranking positions than to fall, because there is a greater opportunity for rising than for falling through the ranking positions.

If it is accepted that the pattern properties of length, overall directional shape, pitch (or duration) centre, and distance between the highest and lowest notes (or longest and shortest notes) can indicate whether PTs are related to each other, then the fact that a number of PTs that first appear in period 3 can be matched with PTs in period 2 suggests the possibility of variation. By then considering the ranking positions of these new period 3 PTs in the subsequent two periods, it can be argued that these indeed show selection, replication and variation.

6.3.2 Exploring Variation in the Music

Investigating the data *en masse* for variation has some drawbacks in terms of providing a suitable measure of distance between different patterns. Chapter 3, Section 3.3 above gives an explanation of the difficulties involved in using computers to determine whether patterns are variants of each other, and discusses the significant problem of trying to provide a suitable algorithm that the computer can implement that is universally applicable to different patterns in different contexts from different periods of time. As a result, this thesis necessarily uses a relatively simple measure by defining PTs as being potential variants if a set of predetermined pattern properties exactly match. However, using statistics based on this measure of variation does not show how effective this approach has been. This section will first examine the viability of this measure of variation (i.e., it will examine the workings of the similarity algorithm) and then assess if this approach can identify suitable variant PTs within the dataset by investigating some PTIs within the scores.

Similarity Algorithm

This section will examine the similarity algorithm by separating out the pitch from the duration properties of the PTs, and investigating the two sets of properties separately. The pitch properties of the PTs will be tackled both using the absolute pitch and the relative intervallic values of the notes without duration (A-PT and R-PT respectively), whilst the duration properties will be tackled by using the relative durations of the notes only without the pitch (there is no corresponding PT descriptor from the four PTDs for duration-only PTs).

To begin, a short three-note A-PT (Example 6.25) is investigated to look at the pitch-only PTs that are deemed to be variants of it according to the pitch properties that are used in the similarity algorithm.

Example 6.25: Three-note PT consisting of the first three degrees of an A major scale

Example 6.25 shows a three-note A-PT that uses the first three notes of the A major scale. Looking at just the pitches (i.e., excluding the duration), the pitch properties for this A-PT are: Notes = 3, Shape = U, PitchCentre = 2, and HighLow = 4 (see Section 4.6.7 for an explanation of these properties).

There are 133 other A-PTs that have the same pitch properties as that in Example 6.25. Because the A-PT refers to actual pitch values, it could be expected that there would be matching pitch properties with all A-PTs that have the first three notes from any of the major scales. However, there are only 12 major keys in total, which means that there are still 122 A-PTs (i.e., 133 other A-PTs with the same pitch properties minus the 11 other major keys) that have matching pitch properties with Example 6.25 that cannot be accounted for using the first three notes of a major scale. Again, because the A-PTs refers to the actual pitch values of the notes, some of these 122 A-PTs that match the pitch properties could be caused by enharmonic equivalents. Example 6.26 is an A-PT that has the same pitch properties as Example 6.25 because it is enharmonically equivalent.

Example 6.26: Enharmonic equivalent to Example 6.25

Those enharmonic equivalent A-PTs within the dataset that match the pitch properties of Example 6.25 contribute another 32 out of the possible 133 matching A-PTs which, taken with the 12 major keys, still leaves 90 A-PTs remaining unaccounted for. The remaining 90 form a group consisting of A-PTs that still begin and end a major third apart, but with the middle note moving either a semitone up or down, two examples of which are shown in Example 6.27.

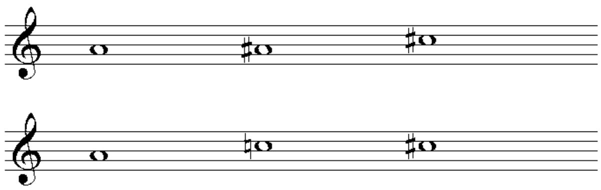

Example 6.27: Two examples of PTs that match the absolute value pitch properties of Example 6.25

Therefore, the majority (90 out of 133) of A-PTs that have an exact match with the pitch properties of Example 6.25 have a different intervallic structure from Example 6.25. Example 6.25 consists of a continuous upward movement of two consecutive whole-tones, whereas 90 of the A-PTs matching the pitch properties of Example 6.25 consist of an upward movement of either a semitone followed by three semitones, or three semitones followed by a semitone, as in the examples in Example 6.27. Under the algorithm used in the present research, both the latter intervallic structures are regarded as having the same pitch properties and are therefore deemed to be potential variants of each other.

When looking at the pitch properties using R-PTs, there are only two other R-PTs that have the same pitch properties as Example 6.25 in its R-PT form (0/U2/U2[13]). The R-PT for Example 6.25 will cover all the pitch-only PTIs that have the same intervallic structure, i.e., two consecutive whole-tone intervals moving upwards. The other two R-PTs that have the same pitch properties as Example 6.25 include the pitch-only PTs in Example 6.27(0/U1/U3 and 0/U3/U1), together with all the A-PTs that have the same intervallic structure as these two examples regardless of their pitch content.

For both the A-PTs and R-PTs, there are some PTs that do not have an exact match of pitch properties as any other PTs within their PTD, examples of which are shown in Example 6.28 and Example 6.29.

[13] See Section 4.4.4 for an explanation of these encodings

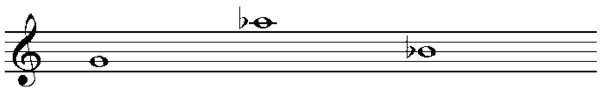

Example 6.28: Three-note A-PT that has no matching pitch properties with any other A-PT

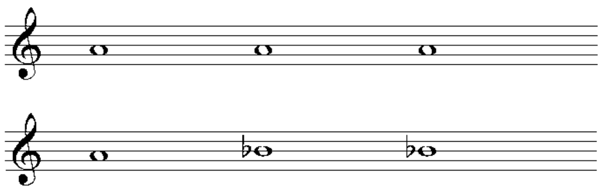

Example 6.29: Two three-note R-PTs that have no matching pitch properties with any other R-PT

All of the three-note A-PTs that do not have an exact match with any other A-PT's pitch properties have a large intervallic distance (of over an octave) between two consecutive notes. However, the same is not true for R-PTs. There are some three-note R-PTs that do not have an exact match with any other R-PT's pitch properties that have intervallic distances that encompass less than an octave between the pitches (see Example 6.29).

When looking at seven-note pitch-only PTs (A-PTs and R-PTs), there is a larger number of PTs that have the same pitch properties as other PTs than is the case for the three-note pitch-only PTs, regardless of whether using the A-PTs or R-PTs descriptor. Because of the way the pitch properties are calculated, there are many more possible pitch-only PTs with the same pitch properties for seven- than for three-note pitch-only PTs. For example, a seven-note R-PT that consists of six ascending whole tones (i.e., 0/U2/U2/U2/U2/U2/U2) can be altered by changing the first interval to a semitone and changing any of the subsequent intervals to a minor 3rd (i.e., 0/U1/U3/U2/U2/U2/U2 or 0/U1/U2/U3/U2/U2/U2, etc.) and still retain the same pitch properties. However, a three-note R-PT that consists of two ascending whole tones (i.e., 0/U2/U2) can only make one such change in order to retain the same pitch properties.

The same logic regarding the increased likelihood of pitch property matches between pitch-only PTs holds true when comparing three- and seven-note PTs to eleven-note PTs. Table 6.20 shows the number of A-PTs and R-PTs that have the same pitch properties as those pitch-only PTs that consists of the first three, seven, or eleven notes of a continuous A major scale starting on the tonic (see Example 6.30).

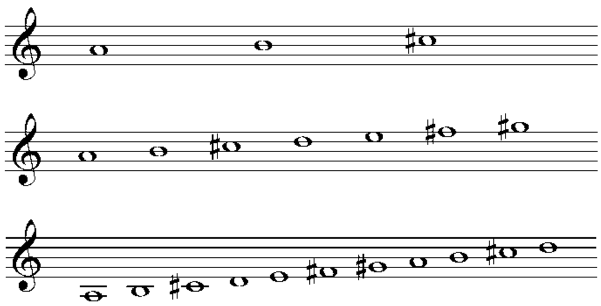

Example 6.30: Three-, seven-, or eleven-note pitch-only PTs consisting of a continuous A major scale starting on the tonic

PTL	PTD	Number of Other Pitch-Only PTs With the Same Pitch Properties as Example 6.30
3	A-PT	134
	R-PT	3
7	A-PT	289
	R-PT	149
11	A-PT	211
	R-PT	90

Table 6.20: Number of pitch only PTs with the same pitch properties as a PT consisting of the first three-, seven- or eleven notes of a continuous A major scale starting on the tonic

Table 6.20 shows that there are more pitch-only PTs with the same pitch properties for the seven- than the three-note PTs from Example 6.30. However, Table 6.20 also shows that there are fewer pitch-only PTs with the same pitch properties for the eleven- than the seven-note PT from Example 6.30, which is contrary to what should be expected according to how the pitch properties are calculated. This is because there is only a larger *potential* number of pitch-only PTs that can have the same pitch properties for the longer pitch PTs. It should be remembered that just because there is a larger number of *potential* pitch-only PTs with the same pitch properties for eleven- than for seven-note PTs, it does not necessarily follow that composers will actually use all the possible pitch-only PTs that can be generated from eleven notes. In other words, there are more possible pitch-only PTs that will have the same pitch properties as those in Example 6.30 than exist in the dataset.

So far, only the pitch properties of PTs have been investigated, without any regard to the duration of the pitches involved. The next stage is, conversely, to look at the similarity of durations without any regard to the pitches. Example 6.31 shows a three-note duration-only PT, and Example 6.32 shows a possible variant that has the same duration properties as Example 6.31, i.e., Notes = 3, Shape = U, DurationCentre = 2.3333 , and HighLow = 7 (see Section 4.6.7 above for an explanation of these properties).

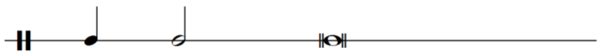

Example 6.31: Sample three-note rhythm

Example 6.32: Sample three-note rhythm with the same duration properties as Example 6.31

The difference between Example 6.31 and Example 6.32, obviously, lies in the fact that the second note in Example 6.32 is twice the length of the second note in Example 6.31. Other than that, both the outer notes in the examples have the same duration. Therefore, it could be argued that the examples are similar, in that the second note is longer than the first, the third is longer than the second, and that the outer notes are the same duration, in both examples.

There are, however, relatively few three-note duration-only PTs that share the same duration properties as any other three-note duration-only PT (82 out of a possible 1,769). When looking at the properties of three-note pitch-only PTs, it was noted that there were far fewer PTs with the same pitch properties when using the R-PT compared with the A-PT descriptor. The same explanation regarding the difference between using relative and absolute values could explain why there are relatively few three-note duration-only PTs that have the same duration properties, because the durations are always calculated using relative rather than absolute values.

When looking at the properties of the seven- and eleven-note duration-only PTs, it would be expected that there would be a greater number of PTs with the same duration properties than for the three-note PTs, because of there being a greater number of possible matches (the same logic as used for pitch properties above). Table 6.21 shows the number of duration-only property matches for the PT in Example 6.33, using the first three, first seven, or all eleven notes of the example.

Example 6.33: Sample eleven-note rhythm

PTL	Number of Other Duration-Only PTs With the Same Duration Properties as Example 6.33
3	1
7	10
11	6

Table 6.21: Number of duration-only PTs with the same duration properties as the first three, seven, and all eleven notes of Example 6.33

The first three notes of Example 6.33 were used above in Example 6.31 to show a possible link with Example 6.32 based on both examples having the same duration properties. Table 6.21 shows that Example 6.31 and Example 6.32 are in fact the only two PTs in the dataset that have the same duration properties as the first three notes of Example 6.33. The first seven notes of Example 6.33 show an increase in the number of duration-only PTs that have the same duration properties, compared to the first three notes. However, the first eleven notes show a *decrease* in the number of duration-only PTs that have the same duration properties, compared to the first seven notes. This is contrary to expectations, because when considering the logic regarding the number of possible PTs sharing the same pitch properties explained above (p. 195), there should be an increase in the number of duration-only PTs with the same duration properties for each subsequently longer PT. Table 6.22 expands on this by giving the total number of duration-only PTs in the dataset, together with the number that have the same duration properties as another duration-only PT.

PTL	Total Nº of Duration-Only PTs	Nº With Same Duration Properties as at Least One Other Duration-Only PT	Percentage With Same Duration Properties as at Least One Other Duration-Only PT
3	1,769	82	5%
7	62,406	16,036	26%
11	195,217	71,754	37%

Table 6.22: Number of duration-only PTs with the same duration properties compared to total PTs

Table 6.22 shows that there are far more seven-note duration-only PTs with the same duration properties than three-note duration-only PTs, both in terms of numbers and as a percentage of the total number of duration-only PTs. However, while there is a substantial increase in the number of duration-only PTs when comparing the seven- to the eleven-note PTL, this is not matched by a corresponding increase in the percentage figure.

There is an additional factor that affects the number of possible duration-only PTs which match the duration properties of other PTs. This is a consequence of how the system measures the duration properties of PTs where all the notes have the same duration. As with one of the examples of the same-note relative pitch from Example 6.29 (three consecutive As at the same octave) that has no other R-PTs with the same pitch properties, the duration-only PT with the same duration across all its notes does not have any other duration-only PTs with the same duration properties.

There is, therefore, an increase in the number of PTs with the same properties from three- to seven-note PTs and then either an increase for duration-only properties or a drop in numbers for pitch-only properties when comparing the seven- to the eleven-note PTs. As explained above, the pitch-only property figures are counterintuitive when considering the possible numbers of PTs having the same properties.

Using both the pitch-only and duration-only properties of PTs to decide if any are variants creates a number of potential matches across all the different PTLs. Some of these matches may not be variants due to the *catch-all* nature of the similarity algorithm. The similarity algorithm also produces a number of PTs that do not have any possible variants based on their pitch-only or duration-only properties. In other words, according to the similarity algorithm, there are some PTs that have not produced any potential variants.

Although the similarity algorithm does highlight certain pitch-only and duration-only PTs as being potential variants of each other, there are some drawbacks to the approach. For example, using an exact match between the pattern properties does not allow for PTs of different lengths to be matched (i.e., variants which add additional notes to a PT). Moreover, the algorithm will not match patterns clearly related by certain compositional techniques, such as inversion.

There is obviously a difference between the A-PTs and R-PTs when using the approach of matching pitch properties to show possible variants of pitch-only PTs. The A-PTs shows all possible A-PTs which have exactly the same relative intervallic structures as each other as being possible variants. By contrast, the R-PTs use the relative distance between the pitches and consequently, all the patterns with the same relative intervallic structure are treated as one PT.

Because durations within the PTs are all calculated on a relative basis (i.e., the duration of each note is expressed as a multiple of the duration of the previous note), the similarity algorithm does not show up matches between duration-only PTs of analogous note duration values. For example, three crotchets will not be shown as being a possible variant of three quavers even though they both have the same relative durations, i.e., they are the same duration-only PT. However, the algorithm still does produce possible variants of duration-only PTs based on their duration properties, as shown in Table 6.22 above. Like the pitch properties, some of the possible variant duration-only PTs may not be true variants due to the *catch-all* nature of the system.

From this point forward, the duration element of PTs will no longer be separated from the pitch element. The music examples will revert to using all four PTDs.

Music Examples

The examples of variant PTs given above have only shown that the similarity algorithm produces some results. The next stage is to determine whether variation actually takes place within the music represented by the dataset. In order to do this, selected new pattern types from period 3 (i.e., those without a PTI in periods 1 and 2) are investigated. These pattern types first appear in period 3, subsequently rise through the ranking positions in periods 4 and 5, and have the highest ranking positions in period 5. Table 6.23 shows these twelve selected PTs together with their ranking positions in periods 3, 4 and 5, and Example 6.34 shows them in standard music notation, with the relative intervallic value PTs starting on c^2, and all durations starting with a crotchet. Example 6.35 then shows PTIs of the PTs in Example 6.34.

PTL	Example Number	PTD	PT	Ranking Position in Period		
				3	4	5
3	3.40 i	A-PT	F/U0F+1/D0F	1,204	905	304
	3.40 ii	AD-PT	E-1:1/D0C-1:1/U0E-1:1	1,004	767	259
	3.40 iii	R-PT	0/U6/U5	530	432	287
	3.40 iv	RD-PT	0:1/D7:1/D9:1	944	535	282
7	3.40 v	A-PT	A/U0B-1/D0A/U0B-1/U0C/D0B-1/U0C	278	141	62
	3.40 vi	AD-PT	G:1/U0A:1/D0G:1/U0A:1/U0B-1:1 /D0A:1/U0B-1:1	195	98	33
	3.40 vii	R-PT	0/U1/U2/R0/D2/S0/U2	632	396	160
	3.40 viii	RD-PT	0:1/D1:1/D1:1/U1:1/U1:1/U1:1/U1:1	468	248	79
11	3.40 ix	A-PT	E/S0E/S0E/S0E/S0E/S0E/S0E/S0E/D0D+1 /U0E/S0E	91	69	34
	3.40 x	AD-PT	D:1/S0D:1/S0D:2/S0D:0.5/S0D:1/S0D:2 /S0D:0.5/S0D:1/S0D:2/S0D:0.5/S0D:1	64	52	29
	3.40 xi	R-PT	0/U4/R0/D4/U4/R0/D4/U4/R0/D4/U4	243	175	76
	3.40 xii	RD-PT	0:1/D1:1/D2:1/U2:1/D2:1/D2:1/D1:1 /U1:1/D1:1/D2:1/D1:1	203	118	36

Table 6.23: Samples of PTs new in period 3 and their ranking positions in periods 3, 4 and 5

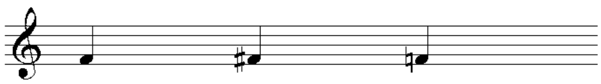

i) A-PT - F/U0F+1/D0F

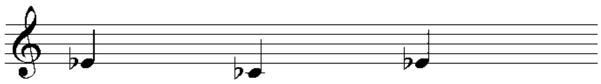

ii) AD-PT - E-1:1/D0C-1:1/U0E-1:1

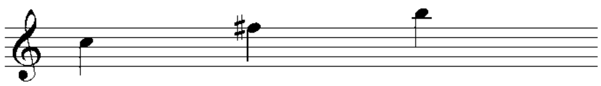

iii) R-PT - 0/U6/U5

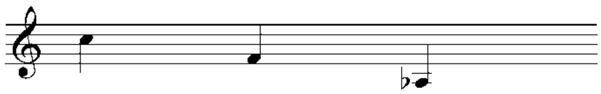

iv) RD-PT - 0:1/D7:1/D9:1

v) A-PT - A/U0B-1/D0A/U0B-1/U0C/D0B-1/U0C

vi) AD-PT - G:1/U0A:1/D0G:1/U0A:1/U0B-1:1/D0A:1/U0B-1:1

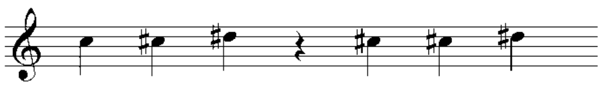

vii) R-PT - 0/U1/U2/R0/D2/S0/U2

viii) RD-PT - 0:1/D1:1/D1:1/U1:1/U1:1/U1:1/U1:1

ix) A-PT - E/S0E/S0E/S0E/S0E/S0E/S0E/S0E/D0D+1/U0E/S0E

x) AD-PT - D:1/S0D:1/S0D:2/S0D:0.5/S0D:1/S0D:2/S0D:0.5/S0D:1/S0D:2/S0D:0.5/S0D:1

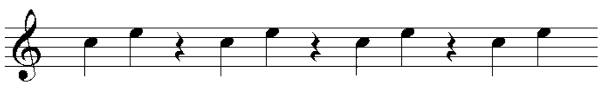

xi) R-PT - 0/U4/R0/D4/U4/R0/D4/U4/R0/D4/U4

xii) RD-PT - 0:1/D1:1/D2:1/U2:1/D2:1/D2:1/D1:1/U1:1/D1:1/D2:1/D1:1

Example 6.34: Examples of the PTs in Table 6.23

i) Beethoven String Quartet op. 74 1st movement bb. 66-68: Instance of PT from Example 6.34 i

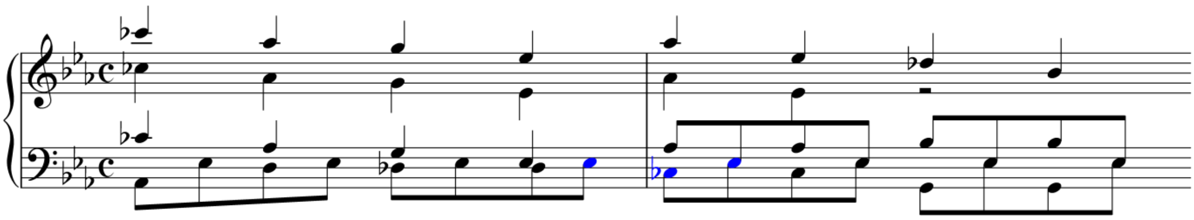

ii) Mendelssohn String Quartet no. 1 1st movement bb. 209-210: Instance of PT from Example 6.34 ii

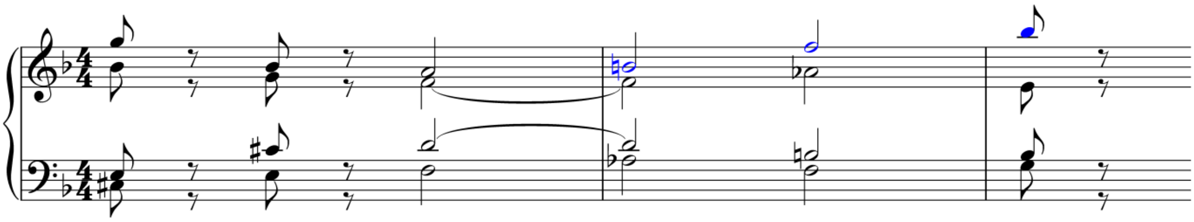

iii) Beethoven String Quartet op. 59 no. 1 1st movement bb. 24-26: Instance of PT from Example 6.34 iii

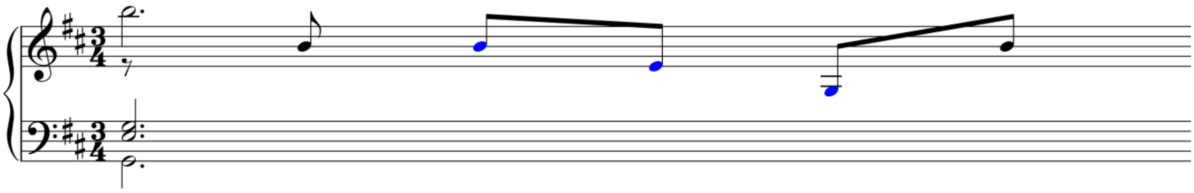

iv) Beethoven String Quartet op. 95 3rd movement b. 70: Instance of PT from Example 6.34 iv

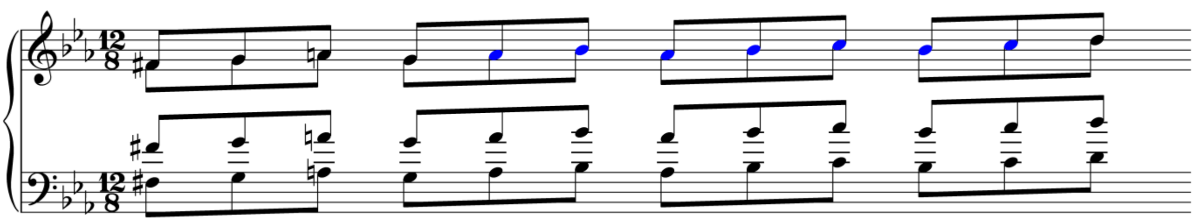

v) Mendelssohn String Quartet no. 1 4th movement b. 95: Instance of PT from Example 6.34 v

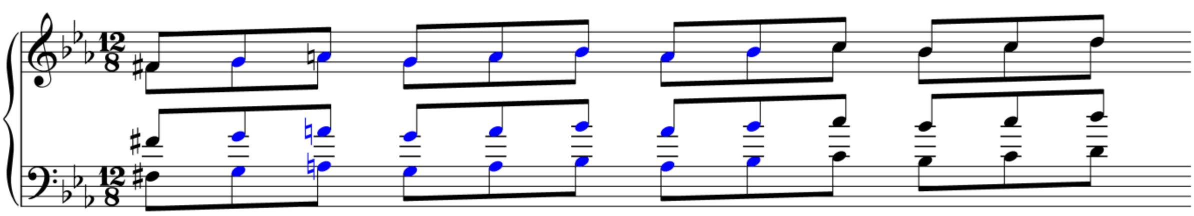

vi) Mendelssohn String Quartet no. 1 4th movement b. 95: Instance of PT from Example 6.34 vi

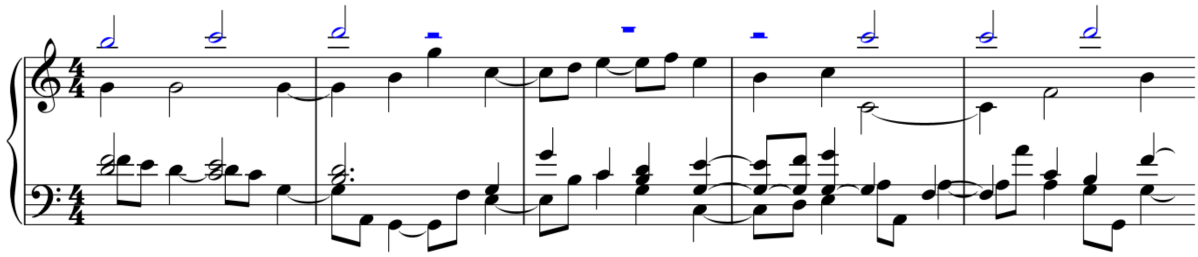

vii) Beethoven String Quartet op. 132 3rd movement bb. 102-106: Instance of PT from Example 6.34 vii

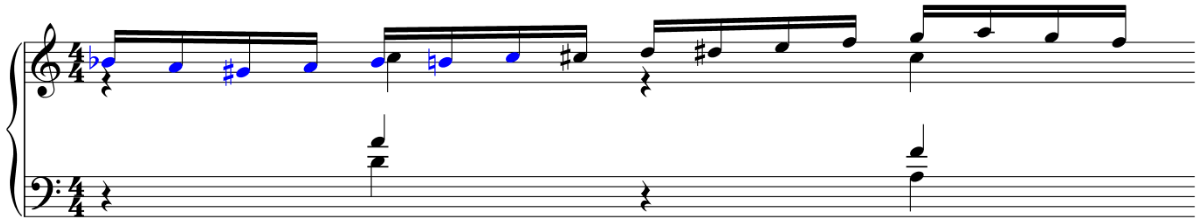

viii) Beethoven String Quartet op. 132 1st movement b. 37: Instance of PT from Example 6.34 viii

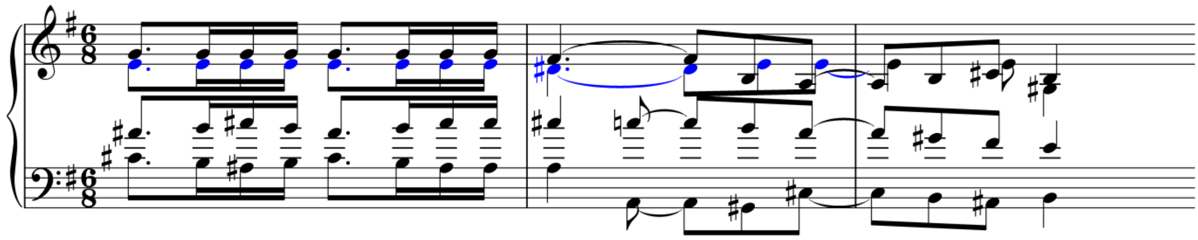

ix) Beethoven String Quartet op. 59 no. 2 1st movement bb. 190-192: Instance of pattern type from Example 6.34 ix

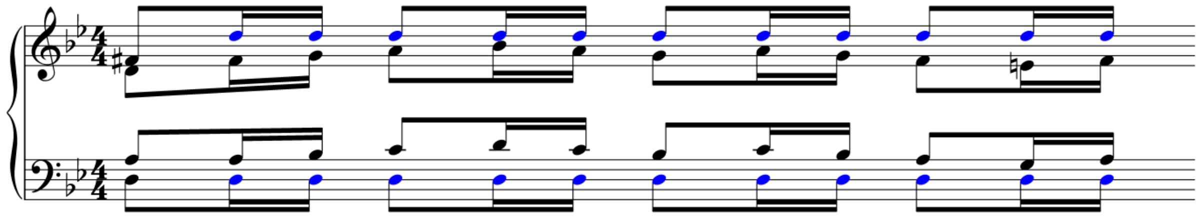

x) Schubert String Quartet no. 14 2nd movement b. 75: Instance of PT from Example 6.34 x

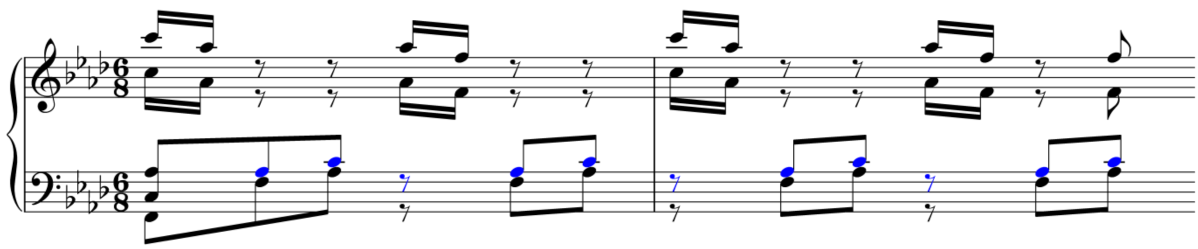

xi) Beethoven String Quartet op. 95 4th movement bb. 11-12: Instance of PT from Example 6.34 xi

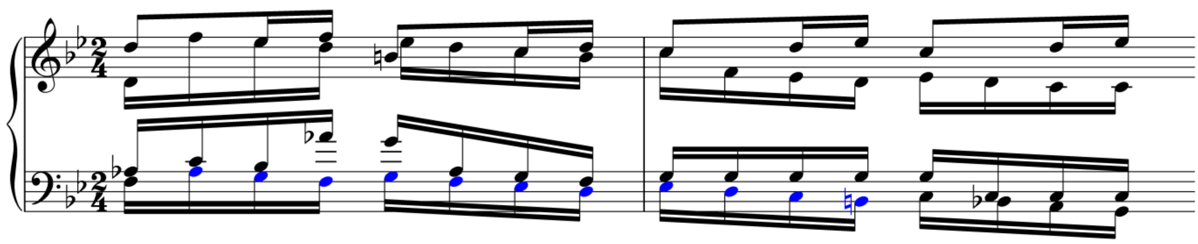

xii) Beethoven String Quartet op. 130 6th movement bb. 207-208: Instance of PT from Example 6.34 xii

Example 6.35: PTIs of the PTs in Example 6.34

The PTs selected in Example 6.34 have a number of similarities in terms of movement, rhythm and tonality mapping with those selected for Section 6.2.2 (Example 6.1 and Example 6.21). For example, most of the PTs in Example 6.34, like most of those in Example 6.1 and Example 6.21, have stepwise movement between the notes. Only the three-note AD-PT, R-PT and RD-PT and the eleven-note R-PT (Example 6.34 ii, iii, iv, and xi respectively) have intervals greater than a major second. In Example 6.1 and Example 6.21 there are six PTs out of a total of eighteen with intervals greater than a major second, of which four are three-note PTs.

It is not just the movement of the pitches in the PTs in Example 6.34 that show a similarity with those in Example 6.1 and Example 6.21. The majority of the PTs in the examples have a very simple rhythmic element, with most of the examples staying on the same length of note. In Example 6.34 only one out of the six examples with a rhythmic element (AD-PT or RD-PT) does not have notes of all the same duration, namely Example 6.34 x. Even this example is not especially complex, in that it has a simple repeated three-note rhythmic pattern. There are also similarities between Example 6.1, Example 6.21 and Example 6.34 when investigating the tonality mappings of the PTs selected, in that most of the PTs in the examples can have their pitches mapped on to scale degrees of the traditional major-minor tonalities.

There are also some similarities between some of the PTs within Example 6.34. For example, Example 6.34 v and Example 6.34 vi follow the same overall shape. The intervallic difference between those two PTs is that whereas one has a semitone interval the other has a whole-tone interval and vice versa. This is also highlighted by the system in that both of the PTs have the same pitch properties: i.e., Shape = UDUDU, PitchCentre = 0.5 and HighLow = 3.

A final point to make is that although the PTs in Table 6.23 progress upwards through the ranking positions across periods 3, 4, and 5, none progresses very far, with no PT reaching a ranking position greater than 29th. When taking into account the number of possible ranking positions for each PTD/PTL/period combination, the eleven-note AD-PT (Example 6.34 x) reaches only the 29th ranking position in period 5, out of a total of 31 possible ranking positions. Table 6.24 shows the ranking position in period 5 of the PTs in Table 6.23, together with the number of ranking positions available for each PTD/PTL combination.

PTL	Example Number	PTD	Ranking Position in Period 5	Nº of Available Ranking Positions in Period 5
3	3.40 i	A-PT	304	511
	3.40 ii	AD-PT	259	394
	3.40 iii	R-PT	287	367
	3.40 iv	RD-PT	282	439
7	3.40 v	A-PT	62	111
	3.40 vi	AD-PT	33	78
	3.40 vii	R-PT	160	234
	3.40 viii	RD-PT	79	162
11	3.40 ix	A-PT	34	44
	3.40 x	AD-PT	29	31
	3.40 xi	R-PT	76	93
	3.40 xii	RD-PT	36	62

Table 6.24: Ranking positions in the last period of the PTs in Example 6.34 together with the number of available ranking positions

Table 6.24 shows that only two out of the twelve PTs in Example 6.34 (vi and viii) make it more than halfway up the available ranking positions in period 5. These are both seven-note PTs that include duration (AD-PT and RD-PT).

The fact that most of the PTs in Example 6.34 do not make the top half of the ranking positions in period 5 suggests that those PTs that first make an appearance in period 3 are not making a significant impact on the ranking positions of the last two periods. In other words, they are not appearing much more frequently in the compositions of the later periods than in period 3, in relation to other PTs.

So far, the PTs in Example 6.34 have been shown to rise through the ranking positions across periods 3 to 5 showing some supporting evidence for selection and replication (as argued in Section 6.2.1 above). These PTs need to be investigated further to see if there are possible PTs in periods 1 and 2 that could have spawned them. Table 6.25 shows the number of possible antecedent PTs in period 2 based on the similarity algorithm.

The numbers for the possible antecedents for each PTL/PTD combination shown in Table 6.25 are tiny compared to the overall number of PTs in period 2. As would be expected, there are fewer antecedents for the R-PTs than the A-PTs as well as fewer antecedents when the duration of the notes are taken into account (with the exception of the eleven-note R-PTs and RD-PTs). Table 6.25 nevertheless shows that, under the similarity algorithm, there are possible antecedents in period 2 for the period 3 PTs shown in Example 6.34. Some of these antecedents will now be explored, both to see how they compare with the PTs in Example 6.34 in terms of their pitches and durations (to investigate the effectiveness of the algorithm), and the scenarios in which they occur (to investigate whether there is any similarity between the scenarios).

PTL	Example Number	PTD	Number of Possible Antecedent PTs in Period 2
3	3.40 i	A-PT	28
	3.40 ii	AD-PT	14
	3.40 iii	R-PT	9
	3.40 iv	RD-PT	7
7	3.40 v	A-PT	140
	3.40 vi	AD-PT	58
	3.40 vii	R-PT	37
	3.40 viii	RD-PT	2
11	3.40 ix	A-PT	51
	3.40 x	AD-PT	8
	3.40 xi	R-PT	23
	3.40 xii	RD-PT	38

Table 6.25: Number of possible antecedents in period 2 of the PTs in Example 6.34

To begin, the similarity algorithm for Example 6.34 i (F/U0F+1/D0F) only identifies relative intervallic equivalent PTs as possible antecedents. Example 6.36 shows a possible antecedent of the PT in Example 6.34 i that is an enharmonic equivalent, together with a PTI from Haydn's String Quartet op. 17 no. 1, from period 1, and another PTI from Mozart's String Quartet K. 387 from, period 2.

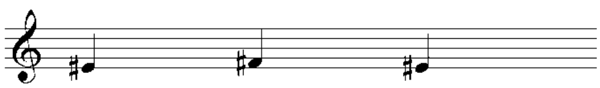

i) Enharmonic equivalent to Example 6.34 i

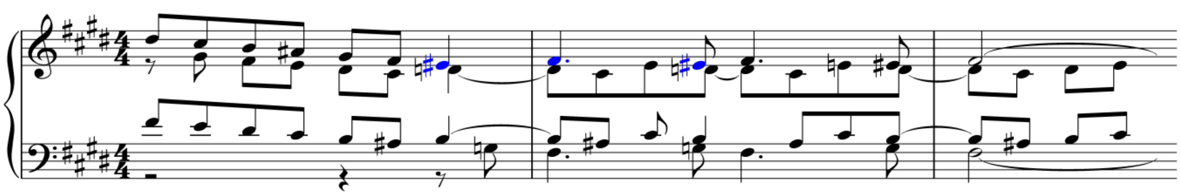

ii) Haydn String Quartet op. 17 no. 1 1st movement bb. 22-24

iii) Mozart String Quartet K. 387 1st movement b. 31

Example 6.36: Possible antecedent of Example 6.34 i that is an enharmonic equivalent

The PTIs in Example 6.36 show these possible antecedent PT to Example 6.34 i in two scenarios. Example 6.36 ii shows the PT as part of melodic interest during the transition from the first to the second subject, whereas in Example 6.36 iii the PT is part of the second subject in the exposition.

Example 6.37 shows the PT in Example 6.34 i transposed up one whole-tone, together with a PTI from Mozart's String Quartet K. 499.

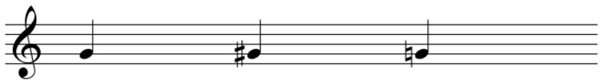

i) A whole-tone transposition of Example 6.34 i

ii) Mozart String Quartet K. 499 3rd movement bb. 93-95

Example 6.37: Possible antecedent of Example 6.34 i that is a whole-tone transposition

There are relatively few PTs (just two within the first two periods in the dataset) which are an exact transposition of Example 6.34 i (i.e., with all three notes having the same pitch name, regardless of the accidentals), with the vast majority being enharmonic equivalents of a transposition (i.e., *without* all three notes having the same pitch name regardless of the accidentals), as in Example 6.38, which shows an enharmonic equivalent transposition together with PTIs in Haydn's String Quartet op. 1 no. 1, from period 1, and from Beethoven's String Quartet op. 18 no. 1, from period 2.

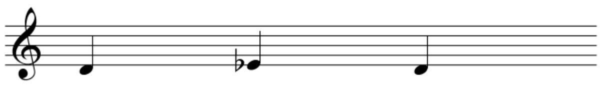

i) A transposition with enharmonic equivalent to Example 6.34 i

ii) Haydn String Quartet op. 1 no. 1 1st movement bb. 0-3

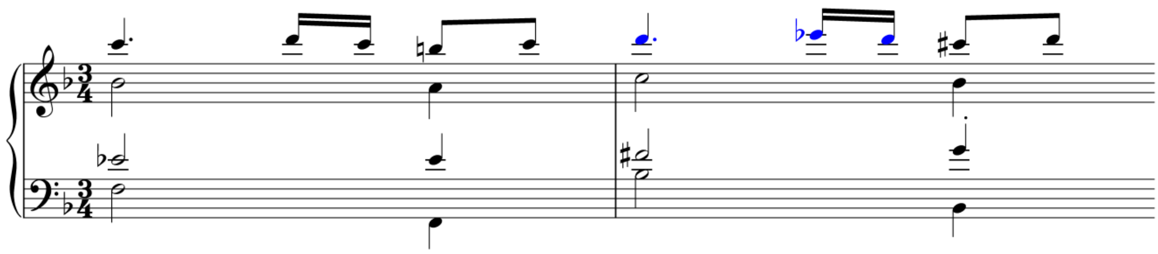

iii) Beethoven String Quartet op. 18 no. 1 1st movement bb. 26-27

Example 6.38: Possible antecedent of Example 6.34 i that have been transposed and use enharmonic equivalents

The possible antecedent (Example 6.38 i) of Example 6.34 i is shown as part of melodic interest in Example 6.38 ii and Example 6.38 iii, which is similar to the PTIs in Example 6.36. With Example 6.38, both PTIs are part of the main thematic material of both the movements shown. Example 6.38 ii occurs at the end of the first phrase of the opening theme, and Example 6.38 iii is in a section that repeats, in a sequence over two bars, the motivic material from the first bar of the movement.

The possible antecedents in Example 6.36 to Example 6.38 reflect the fact that all the possible antecedents in the dataset for Example 6.34 i have the same overall shape and intervallic structure. These examples show that the antecedents are being used as part of melodic material. This is different to Section 6.2.2 above, which showed the selected and replicated PTs in a wider range of scenarios. However, the examples of PTIs of Example 6.34 i do show that they tend to be part of longer patterns, as was the case with the examples showing selection and replication. It must be remembered, though, that this is a small set of samples chosen at random, and therefore, may not be representative of the dataset as a whole.

Moving on to the three-note AD-PT from Example 6.34 ii, both Example 6.34 i and Example 6.34 ii are similar in that they both only produce possible variant PTs that have the same overall shape and intervallic structure. Additionally, for Example 6.34 ii, the duration properties do not produce any possible variants of the rhythmic structure. Example 6.39 shows a possible antecedent PT together with a PTI from Haydn's String Quartet op. 20 no. 4, from period 1, and from Beethoven's String Quartet op. 59 no. 2, from period 2.

i) Enharmonic equivalent to Example 6.34 ii

ii) Haydn String Quartet op. 20 no. 4 1st movement bb. 150-151

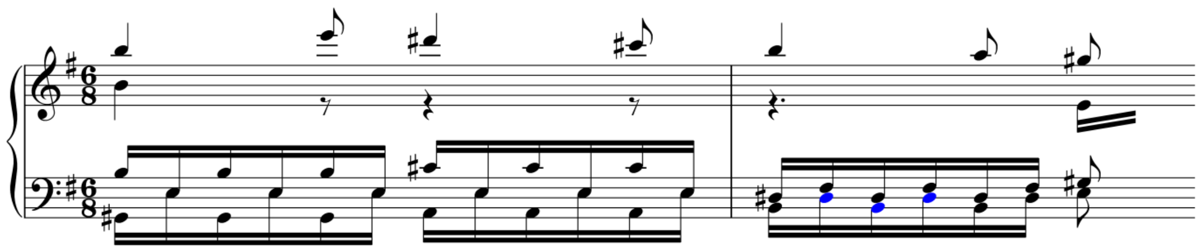

iii) Beethoven String Quartet op. 59 no. 2 1st movement bb. 179-180

Example 6.39: Possible antecedent of Example 6.34 ii that is an enharmonic equivalent

The PTIs in Example 6.39 show the PT in Example 6.34 ii in an enharmonically equivalent version that is used in two different ways. Firstly, the possible antecedent is employed as part of an arpeggiated figure in the first violin part (Example 6.39 ii), and secondly as part of an accompaniment figure in the cello part (Example 6.39 iii).

Example 6.40 shows a possible antecedent to Example 6.34 ii that is a semitone transposition, together with a PTI from Haydn's String Quartet op. 20 no. 3, rom period 1, and from Mozart's String Quartet K. 428, from period 2.

i) Semitone transposition of Example 6.34 ii

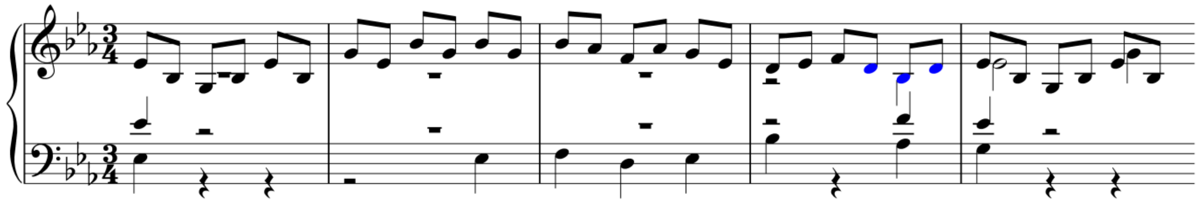

ii) Haydn String Quartet op. 20 no. 3 2nd movement bb. 53-57

iii) Mozart String Quartet K. 428 3rd movement bb. 80-84

Example 6.40: Possible antecedent of Example 6.34 ii that is a semitone transposition

Both of the PTIs in Example 6.40 come from the Minuet and Trio movements. That in Example 6.40 ii begins in the first violin part at the start of the trio section (beginning at bar 53) and is part of the main theme of the section, whilst that in Example 6.40 iii is an accompaniment figure in the second violin part that occurs at the end of the first half of the trio section.

The possible antecedent PTs for Example 6.34 ii in both Example 6.39 and Example 6.40 are, like those for Example 6.34 i, either enharmonic equivalents or transpositions of each other. Unlike the possible antecedents for Example 6.34 i, those for Example 6.34 ii show that the PTIs are being used as part of accompaniment figures as well as part of melodic interest.

When looking at pitches using relative intervallic structure (i.e., the R-PTs and RD-PTs), all patterns with the same intervallic structure are regarded as one PT. Therefore, when looking at R-PTs and RD-PTs, any possible antecedents that the similarity algorithm highlights should have a different intervallic structure from the R-PTs and RD-PTs in Example 6.34 (i.e., iii, iv, vii, viii, xi and xii).

Example 6.41 repeats both the PT and PTI of Example 6.34 iii and Example 6.35 iii respectively (i, ii), along with a possible antecedent with an intervallic change (iii), together with a PTI from Haydn's String Quartet op. 55 no. 1, from period 2 (iv).

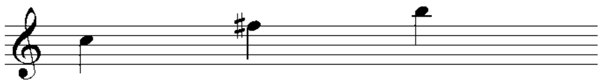

i) PT - Example 6.34 iii

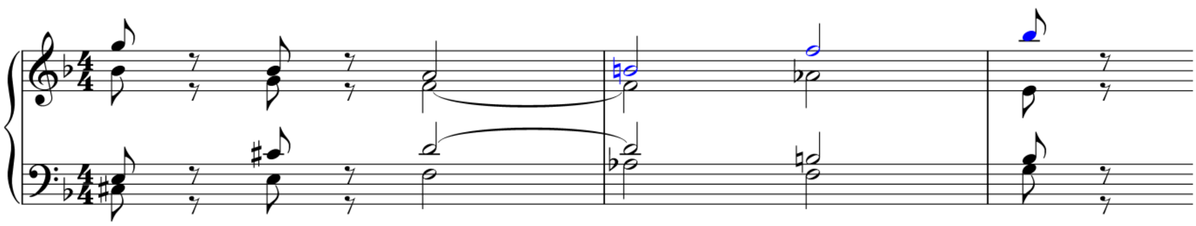

ii) Beethoven String Quartet op. 59 no. 1 1st movement bb. 24-26: Instance of PT from Example 6.34 iii

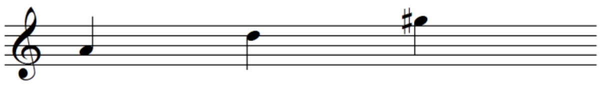

iii) Intervallic change to Example 6.34 iii

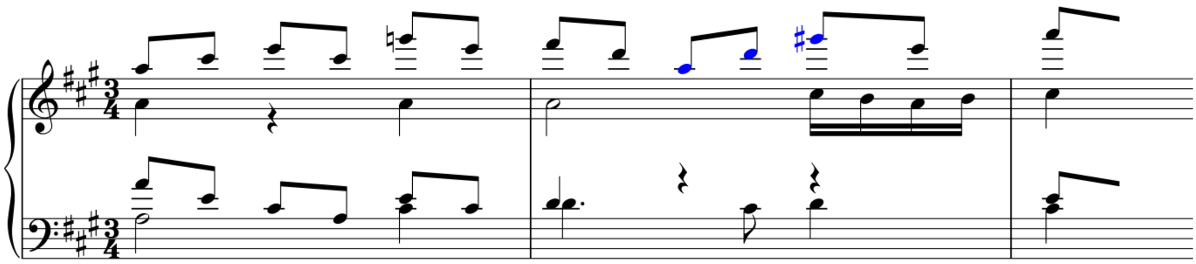

iv) Haydn String Quartet op. 55 no. 1 3rd movement bb. 30-32

Example 6.41: Possible antecedent PT for Example 6.34 iii

In Example 6.41 iii, the possible antecedent to Example 6.34 iii uses the same two internal intervals but in reverse order (i.e., 0/U5/U6 as opposed to 0/U6/U5). Both PTIs in Example 6.41 ii and Example 6.41 iv span an overall interval of 11 semitones, both are in transitional passage where the key of the material is changing, and both have the same internal intervals albeit in a different order.

Interestingly, the PT 0/U5/U6 (Example 6.41 iii) does not have any PTIs in the dataset in period 1. This means that it is new to period 2, and that it may have spawned the PT 0/U6/U5, which is itself new in period 3.

The similarity algorithm also shows other possible antecedents that span the same external interval as Example 6.34 iii but with different internal intervals. Table 6.26 lists all the remaining PTs that are shown by the algorithm as possible antecedents to Example 6.34 iii, together with an indication of whether they appear in period 1 and/or 2.

PT	In Period 1	In Period 2
0/U1/U10	✓	✓
0/U10/U1	✓	✓
0/U2/U9	✓	✓
0/U9/U2	✓	✓
0/U3/U8	✗	✓
0/U8/U3	✗	✓
0/U4/U7	✗	✓
0/U7/U4	✓	✓

Table 6.26: Possible antecedent PTs to Example 6.34 iii

Table 6.26 shows that the majority of possible antecedent PTs to Example 6.34 iii appear in both period 1 and 2 within the dataset. However, three of those in Table 6.26 make their first appearance in the period 2 (as is the case with the possible antecedent PT in Example 6.41 iii). The possible antecedent PTs 0/U3/U8 and 0/U4/U7 are shown in Example 6.42, together with a PTI of each from period 2.

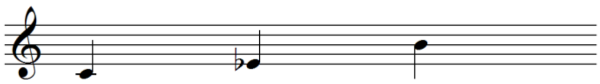

i) Possible antecedent (0/U3/U8) to Example 6.34 iii

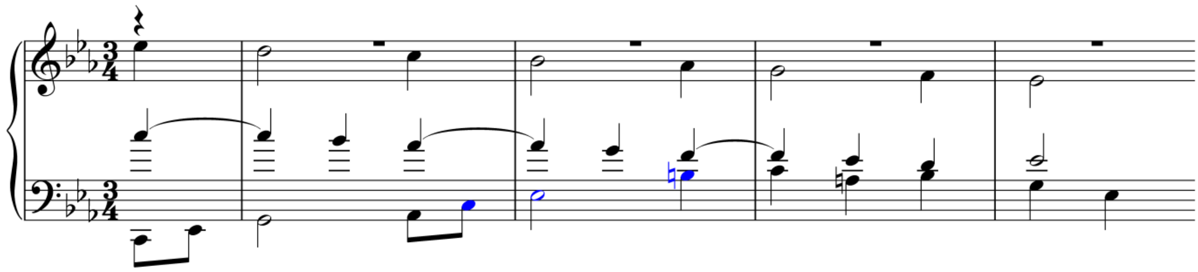

ii) Haydn String Quartet op. 76 no. 6 3rd movement bb. 68-72

iii) Possible antecedent (0/U4/U7) to Example 6.34 iii

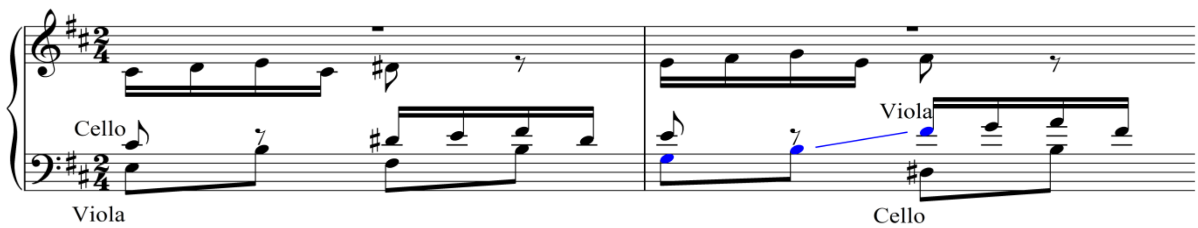

iv) Haydn String Quartet op. 64 no. 2 4th movement bb. 141-142

Example 6.42: Possible antecedents to Example 6.34 iii

The algorithm for identifying possible antecedents has produced some results for Example 6.34 iii which, unlike those for Example 6.34 i and Example 6.34 ii, have various different internal intervallic structures. Additionally, the fact that some of the possible antecedents to Example 6.34 iii make their first appearance in period 2, suggests that they are themselves variants of previous PTs that are then varied again for use in period 3.

Example 6.34 iv shows another set of possible antecedents based on the three-note RD-PT. Example 6.43 repeats the PT and PTI from Example 6.34 iv and Example 6.34 iv, respectively (i, ii), along with a possible antecedent (iii), together with a PTI from Beethoven's String Quartet op. 18 no. 4, from period 2 (iv).

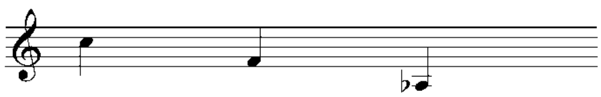

i) PT - Example 6.34 iv

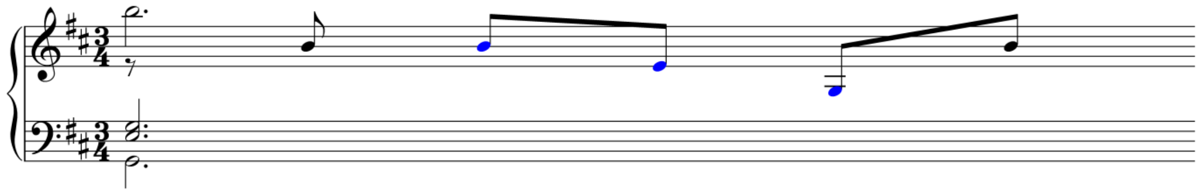

ii) Beethoven String Quartet op. 95 3rd movement b. 70: Instance of PT from Example 6.34 iv

iii) Intervallic change to Example 6.34 iv

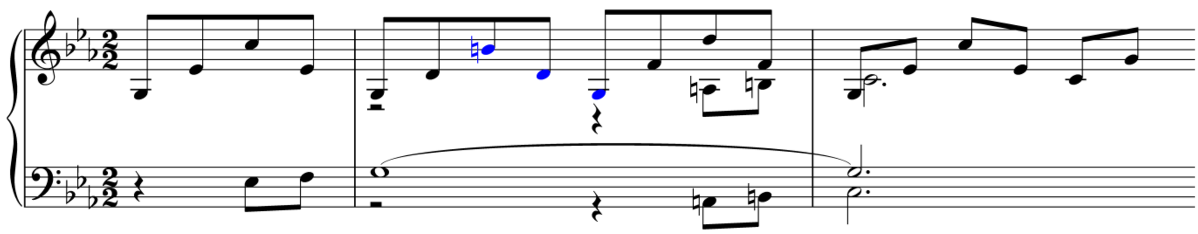

iv) Beethoven String Quartet op. 18 no. 4 4th movement bb. 111-113

Example 6.43: Possible antecedent PT for Example 6.34 iv

Like the possible antecedent in Example 6.41, that in Example 6.43 is based upon the same two intervals as its possible descendant except in reverse order (i.e., 0/D9/D7 instead of 0/D7/D9). However, unlike the possible antecedent in Example 6.41, that in Example 6.43 has at least one PTI in period 1. Example 6.44 shows a PTI of the possible antecedent from Example 6.43 iii in Haydn's String Quartet op. 1 no. 2, from period 1.

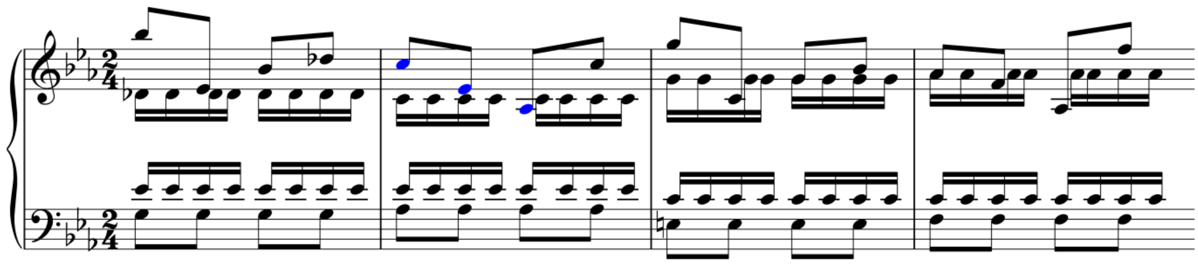

Example 6.44: Haydn String Quartet op. 1 no. 1 5th movement bb. 37-40: Possible antecedent PT for Example 6.34 iv

All the other antecedents to Example 6.34 iv identified in the dataset encompass some of the possible different interval combinations that span a descending major tenth. Table 6.27 lists these, together with an indication whether there is a PTI within period 1 and/or 2.

PT	In Period 1	In Period 2
0/D1/D15	✓	✓
0/D15/D1	✗	✓
0/D2/D14	✗	✓
0/D14/D2	✗	✗
0/D3/D13	✗	✗
0/D13/D3	✗	✗
0/D4/D12	✓	✓
0/D12/D4	✗	✓
0/D5/D11	✗	✗
0/D11/D5	✗	✗
0/D6/D10	✗	✗
0/D10/D6	✗	✗
0/D9/D7	✗	✓
0/D8/D8	✗	✓

Table 6.27: Possible antecedent PTs to Example 6.34 iv

Table 6.27 shows a greater number of possible antecedents than Table 6.26. This is obviously due to the fact that there is a larger interval between the first and last notes in Example 6.34 iv than in Example 6.34 iii, creating a greater scope for more antecedent forms. Another difference between the two tables is that in Table 6.27 there are PTs that do not appear in either period 1 or period 2, whereas all the PTs in Table 6.26 appear in at least one of the first two periods. An explanation for this difference could be that in Table 6.27 (RD-PTs) only those PTs that also matched on the relative duration of the notes were included, whereas in Table 6.26 (R-PTs) the duration of the notes was ignored. Not including the duration properties for the information in Table 6.27, there are only three possible antecedents without any PTIs in the first two periods, compared to seven in the table as it stands.

As explained in the previous section (p. 197), when the duration of all notes in a PT is the same, the similarity algorithm does not identify any possible variants of the durations. Because the PT in Example 6.34 iv has all notes of the same duration (the same is true for Example 6.34 ii), all of the possible antecedents are PTs with notes of the same duration.

Examining the possible antecedents to the three-note PTs in Example 6.34 has highlighted a number of interesting points. These range from differences between the A-PTs and R-PTs arising from how the algorithm identifies possible antecedents, the different PTIs of the possible antecedents and the scenarios in which they appear, and whether the possible antecedents appear in period 1 and/or period 2.

The three-note A-PT and AD-PT (Example 6.34 i and Example 6.34 ii) only supported possible antecedents that had the same underlying intervallic structure, whereas the three-note R-PT and RD-PT (Example 6.34 iii and Example 6.34 iv) produced a range of different underlying intervallic structures. However, the possible antecedents for these two latter types were limited, in that the interval between the first and last note of the PTs was the same, all three notes showed the same direction of progression, and the change was limited to the two internal intervals.

When looking at the possible antecedents to the three-note R-PT and RD-PT, Example 6.34 iii and Example 6.34 iv, some of the possible antecedents only had PTIs in period 2. Therefore, these examples show that there are PTs new in period 3 that have possible antecedents in period 2 that are themselves new to that period.

Because there are more notes in the seven- and eleven- than in the three-note PTs, it would be expected that the longer PTLs would have more possible antecedents than their shorter-length counterparts. This is because there is greater scope for more upward and downward movement within the longer-length PTs, allowing for a greater variety of possible combinations that will share the same pitch properties calculated by the similarity algorithm. What is interesting, however, is that there are fewer possible antecedent PTIs showing up in period 2 for the seven- than for the eleven-note PTs in Example 6.34.

The three-note A-PT and AD-PT patterns in Example 6.34 showed that the only possible antecedents were either enharmonically equivalents or transpositions of the original PTs. However, the dataset does not show any enharmonic equivalent possible antecedents to the seven- or eleven-note A-PTs and AD-PTs in Example 6.34 (v, vi, ix and x). Nevertheless, there are a number of possible antecedents that are transpositions of the seven- and eleven-note A-PTs and AD-PTs in Example 6.34, examples and PTIs of which are shown in Example 6.45.

i) A-PT Example 6.34 v

ii) Possible antecedent PT to Example 6.34 v by transposition

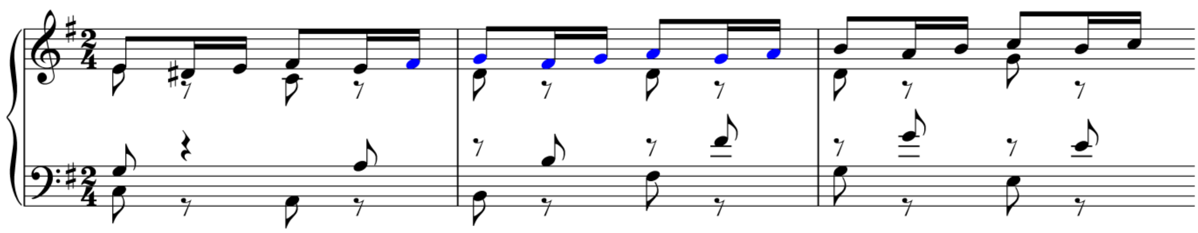

iii) Haydn String Quartet op. 54 no. 1 4th movement bb. 142-144

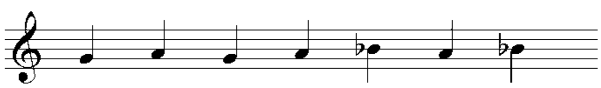

iv) AD-PT Example 6.34 vi

v) Possible antecedent PT to Example 6.34 vi by transposition

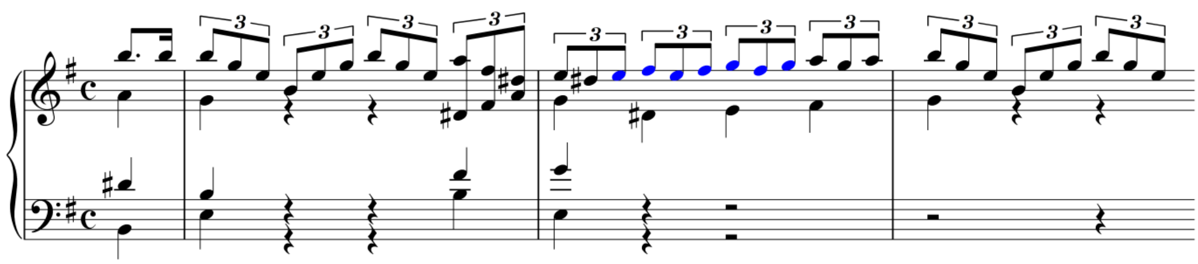

vi) Haydn String Quartet op. 77 no. 1 1st movement bb. 68-71

vii) A-PT Example 6.34 ix

viii) Possible antecedent PT to Example 6.34 ix by transposition

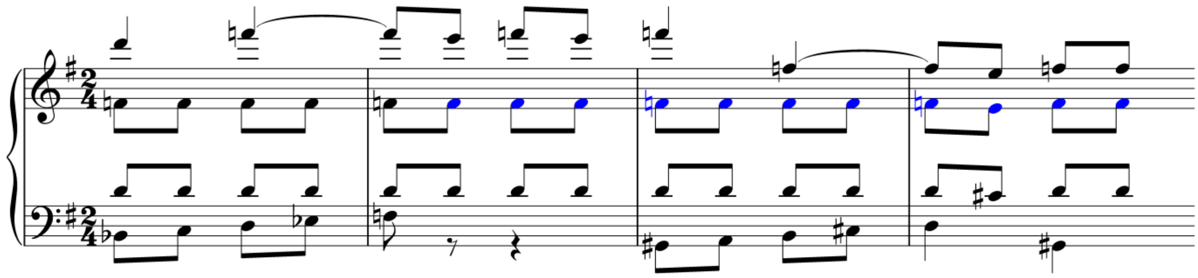

ix) Beethoven String Quartet op. 18 no. 2 4th movement bb. 46-49

x) AD-PT Example 6.34 x

xi) Possible antecedent PT to Example 6.34 x by transposition

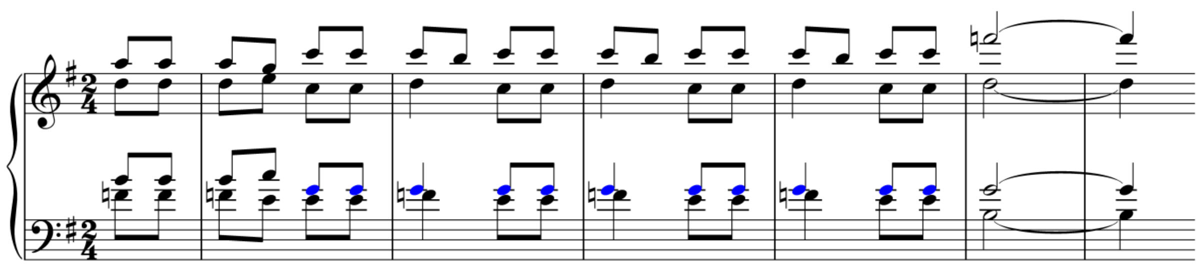

xii) Haydn String Quartet op. 54 no. 2 4th movement bb. 116-122

Example 6.45: Possible antecedents based on transposition of the seven- and eleven-note A-PTs and AD-PTs in Example 6.34

Examples of possible antecedents of the seven- and eleven-note A-PTs and AD-PTs from Example 6.34 that are not direct key transpositions are shown in Example 6.46.

i) A-PT Example 6.34 v

ii) Possible antecedent PT to Example 6.34 v

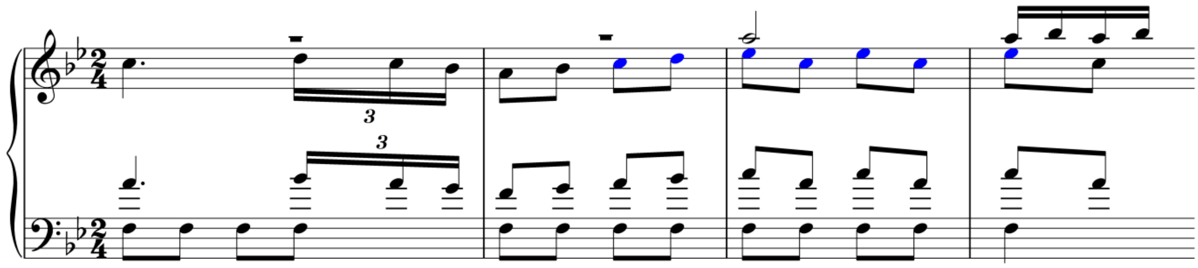

iii) Haydn String Quartet op. 1 no. 1 5th movement bb. 28-31

iv) AD-PT Example 6.34 vi

v) Possible antecedent PT to Example 6.34 vi

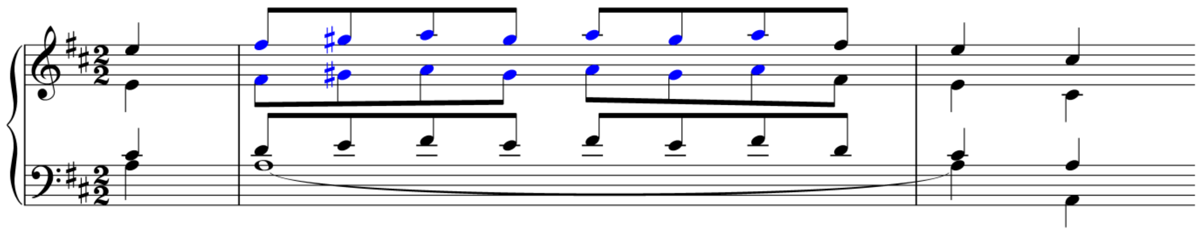

vi) Mozart String Quartet K. 499 1st movement bb. 33-35

vii) A-PT Example 6.34 ix

viii) Possible antecedent PT to Example 6.34 ix

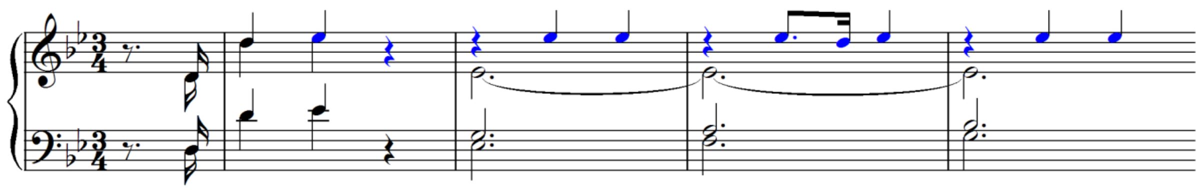

ix) Mozart String Quartet K. 387 2nd movement bb. 105-109

x) AD-PT Example 6.34 x

xi) Possible antecedent PT to Example 6.34 x

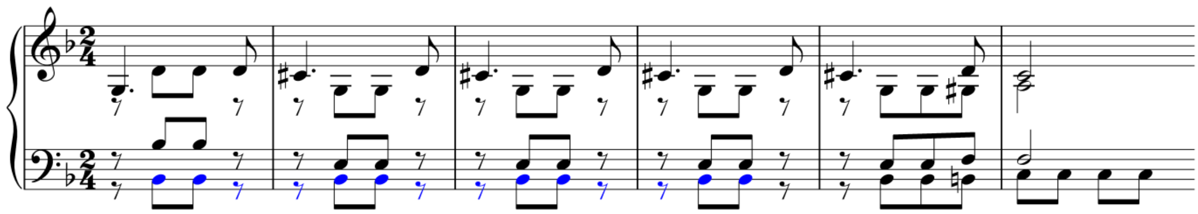

xii) Haydn String Quartet op. 76 no. 2 4th movement bb. 73-78

Example 6.46: Possible antecedents not based on transposition of the seven- and eleven-note A-PTs and AD-PTs in Example 6.34

As discussed earlier in this section (p. 203), both Example 6.34 v and Example 6.34 vi have the same pitch-pattern properties. According to the similarity algorithm, this makes them possible antecedents to each other in terms of pitch, and as such, they both have the same set of antecedents when disregarding the duration of the notes. When taking into account durations, there are fewer possible antecedents to Example 6.34 vi than to Example 6.34 v because all the possible antecedents to the former will consist of PTs that have the same duration for all of the pitches.

The similarity of the pitch element for the possible antecedents to Example 6.34 v and Example 6.34 vi is shown in Example 6.45 ii and Example 6.45 v respectively, where both the possible antecedents share the same direction of the pitches (all the pitches correspondingly rise and fall at the same time) and it is only the intervallic distances that are reversed between the two (i.e., when there is a semitone in one PT, the corresponding interval is a whole-tone in the other and vice versa). However, when considering the other possible antecedents to Example 6.34 v and Example 6.34 vi in Example 6.46 ii and Example 6.46 v respectively, there is a greater distance between the two corresponding PTs in terms of their intervallic content. Whereas the overall shape of the pitches in the PTs remains the same in both Example 6.46 ii and Example 6.46 v, the distance between the notes does not follow the same rule noted above regarding intervallic distances. It could be argued that the difference in the direction of the pitches and the extension of some of the intervals to a minor third in Example 6.46 ii make this PT too distant from the PT in Example 6.34 v to regard the two as possible relatives.

Another possible barrier to claiming that some of the PTs in Example 6.46 are possible antecedents to the respective PTs in Example 6.34 is how rests are dealt with by the system. For example, the PT in Example 6.34 x does not include rests but its corresponding possible antecedent in Example 6.46 xi does. The difference between the two PTs is ultimately down to the rests and these can give a different character to the music, making the possibility of the two PTs being related more remote. Nevertheless, each PT has all the notes on the same pitch, and both follow roughly the same rhythm. Additionally, the rest in Example 6.46 xi always replaces the longer note in Example 6.34 x.

What the PTs in Example 6.45 and Example 6.46 do show, in parallel to the three-note PTs, is that they appear in a range of different scenarios within the compositions. These range from same-note bass-line figures, such as the PTI in Example 6.46 xii, through use as a sequence, such as the PTI in Example 6.45 iii (this PT example lends itself to sequential treatment), to being part of the melodic material, such as the PTI in Example 6.46 vi, and part of a chordal section before a brief reprise of the opening of the movement, such as in Example 6.45 xii.

Example 6.47 to Example 6.50 show possible antecedents to the seven- and eleven-note R-PTs and RD-PTs from Example 40 together with PTIs of the possible antecedents.

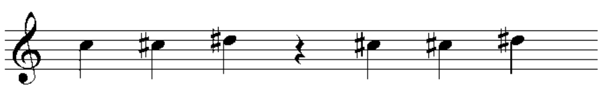

i) R-PT Example 6.34 vii

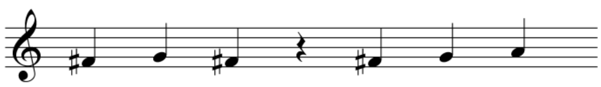

ii) Possible antecedent PT to Example 6.34 vii

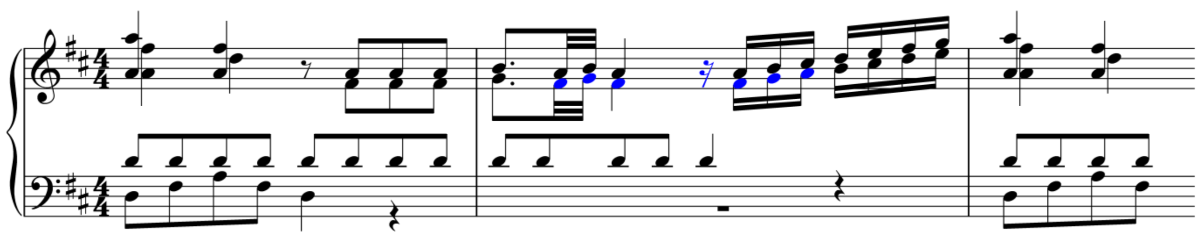

iii) Mozart String Quartet K. 155 1st movement bb. 1-3

iv) Possible antecedent PT to Example 6.34 vii

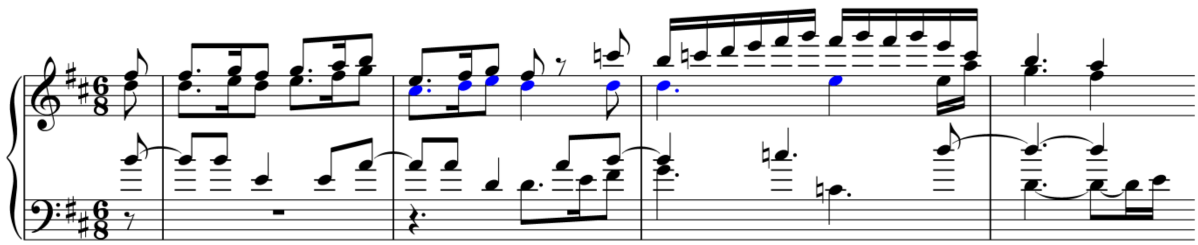

v) Haydn String Quartet op. 33 no. 5 4th movement bb. 41-45

Example 6.47: Possible antecedent PTs to Example 6.34 vii

Both the possible antecedent PTs in Example 6.47 have similarities with their potential descendant PT from Example 6.34 vii. For example, both Example 6.47 ii and Example 6.34 vii have a rest in the same position within the pattern, both have an upward movement spanning three semitones from start to finish, and both have a repeated note (although in Example 6.47 ii the repeated note straddles a rest). However, when looking at the two component three-note patterns either side of the rest independently in each of the PTs in Example 6.34 ii and Example 6.47 vii, there is less similarity between the PTs in terms of the intervals used and the shapes of the patterns. It could be argued that Example 6.34 vii is closer to Example 6.47 iv than to Example 6.47 ii because the only difference between them is that Example 6.47 iv has a note (the fourth note) where Example 6.34 vii has a rest.

i) RD-PT Example 6.34 viii

ii) Possible antecedent PT to Example 6.34 viii

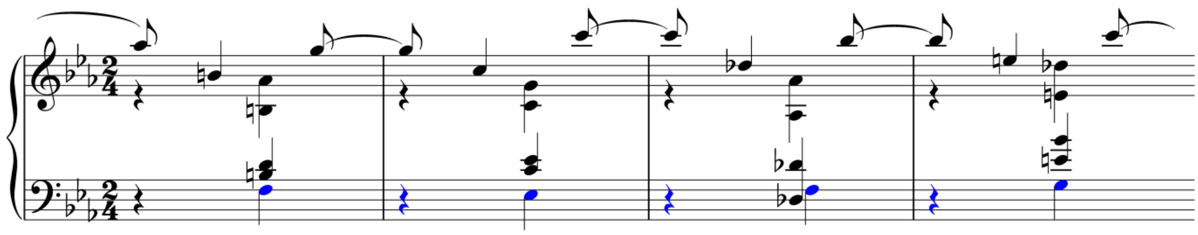

iii) Haydn String Quartet op. 20 no. 1 4th movement bb. 80-83

iv) Possible antecedent PT to Example 6.34 viii

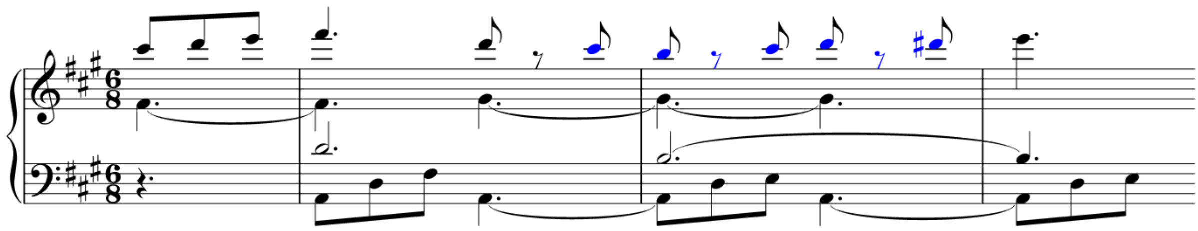

v) Beethoven String Quartet op. 18 no. 5 1st movement bb. 4-7

Example 6.48: Possible antecedent PTs to Example 46 viii

The possible antecedent PTs to Example 6.34 viii in Example 6.48 show different distances from their hypothesised descendant in terms of shape, intervallic structure, and the use of rests. Example 6.48 ii has the same overall shape and the same durational values for the notes as Example 6.34 viii, with even-numbered notes replaced by a rest. All of the intervals between odd-numbered notes in both PTs remain the same. However, although Example 6.48 iv follows roughly the same directional shape and has the same durational values for the notes as its possible descendant, the rests cannot be replaced with any notes that would create the same intervallic structure as its possible descendant, as could be done with the PT in Example 6.48 ii. This creates the odd situation whereby the PT with just four pitches and three rests (Example 6.48 ii) could be considered a stronger candidate for an antecedent to Example 6.34 viii than the PT with five pitches and two rests (Example 6.48 iv).

i) R-PT Example 6.34 xi

ii) Possible antecedent PT to Example 6.34 xi

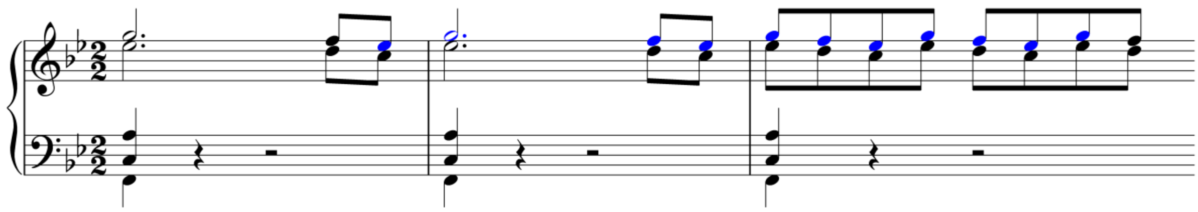

iii) Haydn String Quartet op. 76 no. 4 4th movement bb. 20-22

iv) Possible antecedent PT to Example 6.34 xi

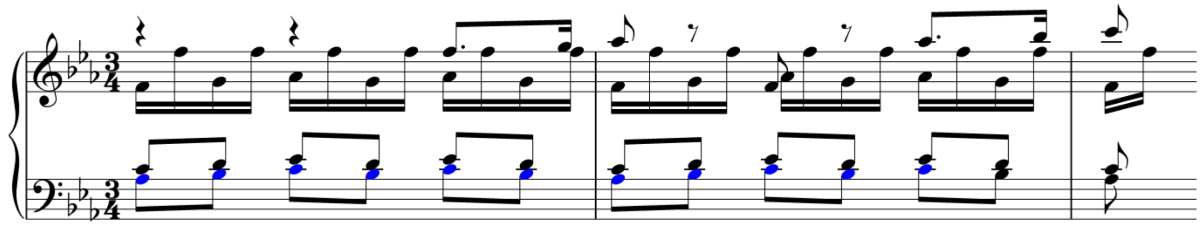

v) Mozart String Quartet K. 589 3rd movement bb. 86-88

Example 6.49: Possible antecedent PTs to Example 46 xi

At first glance, Example 6.49 ii looks quite different from Example 6.34 xi. However, on closer inspection the two could be considered similar if the rests were replaced in Example 6.34 xi with a passing note between the E and C; then both PTs would have exactly the same intervallic structure. Example 6.49 iv, however, looks completely different from Example 6.34 xi and there appears little discernible connection between them. For example, Example 6.34 xi has a repeated three-note pattern with a rest on the third beat, whereas Example 6.49 iv consists of the same five-note pattern (the first and last five notes) separated by a further note without any rests. Therefore, it is questionable whether Example 6.49 iv could be a possible antecedent to Example 6.34 xi.

i) RD-PT Example 6.34 xii

ii) Possible antecedent PT to Example 6.34 xii

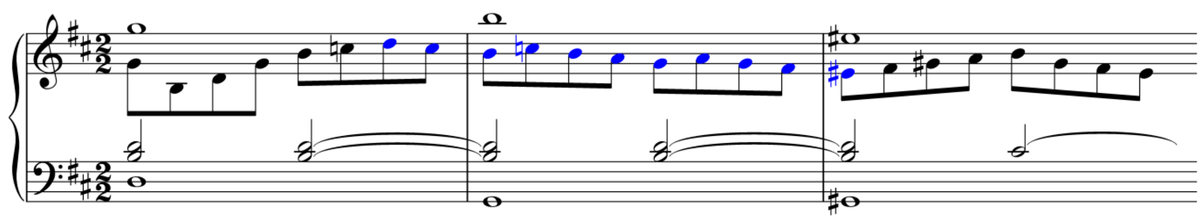

iii) Beethoven String Quartet op. 18 no. 3 1st movement bb. 144-146

iv) Possible antecedent PT to Example 6.34 xii

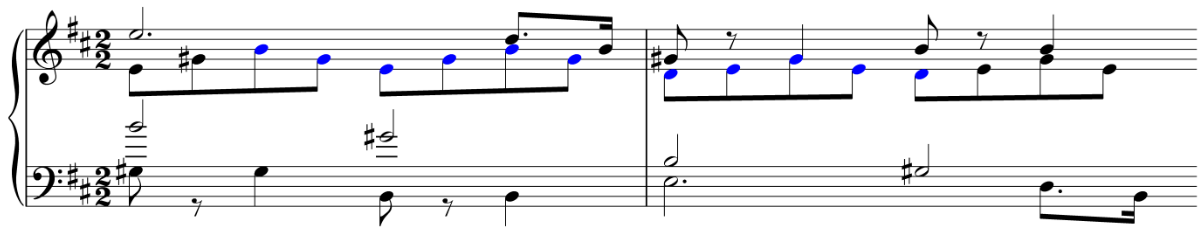

v) Mozart String Quartet K. 499 1st movement bb. 44-45

Example 6.50: Possible antecedent PTs to Example 46 xii

The possible antecedent PTs in Example 6.50 are comparable to those in Example 6.49 in that the first potential antecedent PT in each example matches their respective potential descendants more closely than the second. Example 6.50 ii follows the same shape as Example 6.34 xii (i.e., three consecutive downward notes, then up a step, then four consecutive downward notes, then up a step, then four more consecutive downward notes) as well as having all the notes of the same duration. The difference between the two PTs is that sometimes a semitone is replaced by a whole-tone and vice versa. Example 6.50 iv has a less readily discernible link to Example 6.34 xii, in that although it still has three downward movements with upward movements in between, and all the notes are of equal duration, the intervals are mostly greater than a major 2nd (the largest interval in Example 6.34 xii).

The possible antecedent PTs in Example 6.47 to Example 6.50 show a wide variation in similarity to their respective possible descendants in terms of directional shape, intervallic content, and the use of rests. Some of the possible antecedent PTs bear a close relationship to their possible descendants, such as between Example 6.50 ii and Example 6.34 xii, whereas others seem to bear little, such as between Example 6.49 iv and Example 6.34 xi.

6.3.3 Variation Summary

It was shown that there are PTs within period 3 for all four PTDs that do not appear in either period 1 or period 2. The question that is then raised is whether these PTs new in period 3 are descendants of existing PTs, i.e., those that have appeared in the first two periods. Using the similarity algorithm, it was found that the PTs new in period 3 could be descendants of PTs from period 2. It was also found that some of these PTs new in period 3 showed signs of selection and replication by moving upwards through the ranking positions across periods 4 and 5. However, the figures for the number of PTs new in period 3 that progress upwards through the ranking positions are relatively tiny compared to the total number of new PT's in that period. Additionally, there were no PTs new in period 3 that continuously progressed downwards through the ranking positions, although this could be a consequence of the number of possible ranking positions being smaller in the later than the earlier periods

For variation to have been shown in the data, the similarity algorithm needs to be accepted as an accurate way of underpinning variation between PTs. When looking at the PTIs of selected PTs new in period 3 and their potential antecedents, it was discovered that some of the matches could be considered to be variants of each other, but that others would be considered by musicians as improbable.

Although the data has shown that there are a number of PTs new in period 3 that can be related to PTs from earlier periods using the similarity algorithm, it is difficult to say that the new PTs are potential memes for two reasons. Firstly, the similarity algorithm is not completely accurate as a means of determining a musical variation between PTs; and, secondly, the difference in the number of ranking positions between periods has meant it is difficult to determine whether the PTs in period 3 are progressing upwards or downwards through the ranking positions in subsequent periods and therefore whether the PTs are indeed exhibiting selection and replication.

6.4 Summary

The initial stage of this chapter was to investigate the data for selection and replication, as stated in the working definition of a meme in music. It was argued that looking at the movement of the PTs through the ranking positions across the periods would indicate whether there was selection and replication taking place. This was on the basis that if a PT's frequency of appearances increased over time, then it was showing evidence of being selected and replicated, or of being rejected if its frequency decreased. A number of PTs were found that progressed either continuously upwards or continuously downwards through the ranking positions across consecutive periods, with some of these either not appearing in period 1 or period 5 according to the direction of their progression. The example of the three consecutive Gs A-PT mentioned above shows some movement, in that it starts in ranking position 3 in period 1 and falls to position 81 in period 5. Therefore, this PT is used more frequently in the earlier than in the later periods, compared to the frequency of the other three-note A-PTs in those periods.

Unfortunately, the figures for the number of PTs that continuously progress upwards or downwards through the ranking positions across the periods is very small compared to the total number of PTs in each period. Additionally, the large majority of PTs (over 90% for each PTD) only occur in one of the five periods. Moreover, only a tiny fraction of PTs (fewer than 0.5% for each PTD) appear in all five periods. However, these figures could be a result of the methodology, which generates all possible patterns regardless of any secondary parameters such as metre, and consequently includes many patterns that would not necessarily be identified as such by a composer, performer or listener.

When looking at the PTIs of the PTs sampled, a couple of points were raised. Firstly, the PTIs were, on the whole, part of a longer phrase or motive. Many of the PTIs appeared in the middle of a phrase or motive and began on a weak metric position. Secondly, the PTIs appeared in many different scenarios across all the periods for all of the PTD/PTL combinations. Both of these points suggest that there is little importance to the placement of the PTs within the music when looking at selection and replication using the present methodology. Additionally, it should be remembered that a very small number of the PTIs were investigated.

The next stage was to search for evidence of variation by looking at those PTs that first appeared in period 3 and determining whether they had any antecedents. Each of the new PTs from period 3 investigated showed that it had some potential antecedents. However, when comparing these new PTs to their hypothesised antecedents, it was shown that some antecedents highlighted by the methodology were improbable. Nevertheless, there were a number of antecedents identified that could be described as strong candidates to be possible antecedents. Therefore, there is some evidence for variation based on the similarity algorithm implemented.

The new PTs from period 3 were also investigated for their movement within the ranking tables of the subsequent periods for potential evidence of selection and replication. This found that there were some PTs new in period 3 that did continuously progress upwards through periods 4 and 5. Both their movement through the ranking positions and the possibility of their having antecedents suggests that some of these PTs new in period 3 are indeed exhibiting selection, replication and variation. However, it should be remembered that there are fewer ranking positions available in the later than the earlier periods, increasing the ease of PTs being able to move upwards through the ranking positions.

Unlike the investigations into selection and replication, there was some limited evidence that potential antecedent PTs of the PTs new in period 3 tended to be part of melodic interest. This evidence was restricted to certain three-note A-PT and AD-PTs, with all the other PTL/PTD PTIs used as examples showing a wider variety of scenarios. However, as with the investigation into the scenarios of PTs for selection and replication, the potential antecedent PTIs tended to be part of longer phrases or motives. Again this points to the placement of PTIs within the music having little bearing on memes in the present methodology.

7 Chapter 7: Replicator Properties

7.1 Introduction

Following on from the previous chapter which looks at the evolutionary processes of selection, replication and variation, this chapter looks at the replicator properties of longevity, fecundity and copying fidelity. It is not possible to use the ranking positions as direct evidence for these because they relate more to the individual instances (i.e., one instance in a single printed edition) of a meme rather than the meme itself (i.e., the overall structure of all the instances of the same meme). However, evidence can be inferred from the ranking positions by looking at the relative abundance of patterns. This argument is based on Dawkins' idea that the ultimate aim of a gene, and therefore a meme, is to become dominant within the gene, or meme, pool (1989, pp. 16-17). Therefore, those patterns that show a greater abundance within the dataset will, in order to have gained this greater abundance, have exhibited greater longevity, fecundity and copying-fidelity than those less abundant patterns.

7.2 Evidence for Replicator Properties

The replicator properties of longevity, fecundity and copying-fidelity are more difficult to find evidence for than the evolutionary processes because the properties relate more to the individual manifestations (i.e., instances) of the memes rather than the meme itself (i.e., the overall entity that encompasses all the individual instances of a particular meme). A meme instance can appear in a number of different storage mechanisms (as discussed in Section 2.5.1 above) including the score that is used for the present research. Each of these storage mechanisms will have a different bearing on the three properties. For example, a melody in an autograph score from the eighteenth-century stored in a public library exhibits greater longevity than the same melody stored in the brain of a person who was born in the nineteenth-century. However, if the same autograph score is locked away without the possibility of access, then the memes contained within the score will have little fecundity because there is limited scope for reproduction, whereas the melody stored in the brain of the nineteenth-century person has the chance for high fecundity because the person can recall the meme by methods such as singing or writing it down multiple times. Direct evidence for the replicator properties within the dataset is difficult to determine because the dataset only holds information on one instance (a specific edition) of one storage mechanism (the score) of the potential memes.

Dawkins makes the point that within the pool of DNA molecules, some molecules will survive longer than others, some will reproduce at a faster rate, and some will have copying errors. The result of these processes is that some of the molecules will become more numerous than others, i.e., they will become more abundant within the pool of molecules (1989, pp. 16-17). Using the gene/meme analogy it is possible to argue that the corresponding replicator properties have had an impact on the PTs by looking at their relative frequency within the dataset. For example, those PTs that show a greater abundance within music will, in order to have attained that greater abundance, have exhibited greater longevity, fecundity and copying-fidelity than those less abundant PTs. As such, the three replicator properties are investigated in terms of the relative frequency of PTs.

7.3 Evidence for Longevity

As explained in Section 2.3, longevity relates to the length of time that a meme instance exists. The greater the longevity of a meme instance, in whatever form, the greater the chance for selection and replication to take place and, as a consequence, the more abundant within the meme-pool the meme is likely to become. This abundance then affects the meme's chances for further replication and selection. If a meme's only instances are in a composition that is popular for just a few months, there are fewer chances that it will be noticed by other composers, performers and listeners when compared to a meme with instances in a composition that has lasting popularity across the centuries. When a composition becomes well-known it is obviously more widely performed, studied, recorded, etc. than a little-known composition, providing the meme instances within the well-known composition with a larger potential audience. This audience then has, by extension, more opportunity to select and replicate the memes.

There is, however, a problem in determining what constitutes an instance of a meme. Music can produce a number of different manifestations of memes, such as a score, a compact disc, an mp3 file, a live performance, or even a memory in the brain. All of these manifestations can exhibit the property of longevity, and each has its own different possible time span. For example, an autograph score could have been produced in the eighteenth-century when a composition was created and placed in a public library. This autograph score can then remain in the public domain until the present day, spawning a number of different editions. However, a memory of a meme in the brain could arguably last only as long as it remains in the forefront of consciousness, and at most the recollection can only survive as long as the host survives. Therefore, this recollection in the brain may have fewer opportunities to spawn reproductions than the example used of an autograph score.

In this thesis, the details of the score (i.e., when first published, the number of different publications, etc.) have not been stored in the dataset. It is only the notes that have been taken as the manifestation of memes to be investigated. Therefore it is not possible directly to provide evidence from the dataset for longevity, because the information regarding the availability of the scores across a time period is not part of the dataset.

It may, however, be possible to use the relative abundance of a PT to infer from the dataset evidence for the longevity of memes in music. Because an outcome of longevity is the greater abundance of some memes over others, and because the ranking system used in the present research determines which PTs have been propagated most widely, then it may be possible to use the ranking system to determine which PTs have the greatest abundance and which may consequently have exhibited greatest longevity. Unfortunately, a major drawback to this method is that it ranks PTs based on their prevalence by calculating both the number of times a PT appears in movements as well as the number of different movements in which a PT appears. This results in the possibility that a PT which may appear a great many times in just one or two movements has roughly the same ranking position as a PT appearing only once or twice in many different movements. Another drawback with using the ranking positions to infer longevity is that it is not providing direct evidence for longevity; looking at the ranking positions to infer longevity is ultimately about determining the abundance of certain PTs and assuming that the abundance is a manifestation of longevity. Even though there are drawbacks to inferring longevity from abundance, it is, nevertheless, the only method that is available from the dataset used in this thesis.

The ranking positions are investigated on the assumption that if a PT appears across several periods, then the original PTI may have possessed high longevity in order for it to have had time to be replicated widely. Table 7.1 shows three-, seven- and eleven-note PTs with the greatest continuous progression upwards through the ranking positions across all the periods together with the number of compositions that each PT appears in within each period, with Table 7.2 showing the total number of compositions within each period.

PTL	PTD	PT	Nº and Percentage of Compositions per Period									
			1		2		3		4		5	
3	A-PT	A/D0A-1/D0G	8	32%	23	52%	15	71%	13	100%	7	88%
	AD-PT	A:1/D0G+1:1 /S0G+1:1	7	28%	25	57%	13	62%	11	85%	7	88%
	R-PT	0/D3/U4	7	28%	22	50%	16	76%	11	85%	7	88%
	RD-PT	0:1/D4:1/U3:1	6	24%	23	52%	16	76%	12	92%	6	75%
7	A-PT	E/U0F/D0E/U0F /D0E/U0F/D0E	3	12%	16	36%	9	43%	5	38%	4	50%
	AD-PT	NONE	-	-	-	-	-	-	-	-	-	-
	R-PT	0/U2/U2/U1/U2 /U1/U2	6	24%	20	45%	12	57%	7	54%	4	50%
	RD-PT	0:1/D1:1/D2:1 /D1:1/U1:1/U2:1 /U1:1	5	20%	16	36%	10	48%	5	38%	2	25%
11	A-PT	NONE	-	-	-	-	-	-	-	-	-	-
	AD-PT	NONE	-	-	-	-	-	-	-	-	-	-
	R-PT	0/S0/R0/S0/S0 /R0/S0/S0/R0/S0 /S0	2	8%	16	36%	11	52%	9	69%	5	63%
	RD-PT	0:1/U1:1/D1:1 /U1:1/D1:1/U1:1 /D1:1/U1:1/D1:1 /U1:1/D1:1	2	8%	13	30%	13	62%	7	54%	5	63%

Table 7.1: Number and percentage of compositions in each period for each selected PT

Period	Total Number of Compositions
1	25
2	44
3	21
4	13
5	8
All Periods	111

Table 7.2: Total number of compositions within each period

There are three PTL/PTD combinations that do not have any PTs that continuously progress upwards through the ranking positions across all the periods (the seven-note AD-PT, and the eleven-note A-PT and AD-PT). These PTL/PTD combinations are therefore unable to show any increasing abundance across the periods using this particular method, and it is, therefore, not possible to infer longevity from them as a result.

The PTL/PTD combinations that do have PTs that continuously progress upwards through the ranking positions across the periods show a mixed picture in terms of abundance. Table 7.1 shows that, when compared with the total number of compositions within each period (Table 7.2), for the majority of the selected PTs there is a smaller proportion of compositions containing the selected PTs in period 1 than in period 5. This change across the periods can be more clearly seen when calculating the percentage abundance of the PT, i.e., the number of compositions the PT appears in compared to the total number of compositions within that period. For example, the highest percentage abundance in period 1 is 32%, whereas the highest percentage abundance for period 5 is 88%.

Table 7.1 therefore shows that there are some PTs that do become more widely used through the periods in terms of the percentage of compositions in which they appear. As a result, it could be argued that this spread across the compositions is akin to an increasing abundance of the PTs, especially when compared to other PTs that continuously fall through the ranking positions across the periods (see Table 6.3, on p. 148). If this method for calculating abundance is accepted, then it is possible that longevity exists, because there are some PTs exhibiting greater abundance than others.

Unfortunately, not only are there some PTL/PTD combinations that do not have a selected PT in Table 7.1, there are also PTs in the table that do not show a high degree of abundance in terms of the number of compositions they appear in within period 5. For example, the seven-note RD-PT only appears in two compositions (out of a possible eight) in period 5, which in percentage terms is just 25%. This is not a significant increase from period 1, where the percentage is just 20%.

A further problem with looking at only the note data is that it is impossible to say whether a composer has prior awareness of a particular PT before incorporating that PT into a composition. However, on the whole composers make a conscious effort to familiarise themselves with the works of their contemporaries and previous composers. Therefore, it could be expected that as part of this familiarisation process PTs will be passed between composers. It could also be the case that some of the generated PTs in the system would not be considered as such by musicians. For example, the composer could intentionally create two phrases next to each other without regard for the patterns created across phrase boundaries. However, the pattern generation created patterns across such boundaries. This problem means it is difficult to assess whether an abundance of a PT is due to longevity or is largely coincidental.

The case for longevity is therefore difficult to provide evidence for from within the dataset used for this thesis. However, Dawkins suggests that the individual instances of memes are 'relatively unimportant' (1989, p. 194) for the continued survival of a meme. He explains that the important factor in longevity is that memes need to survive long enough to be passed from generation to generation (1989, p. 194). What Table 7.1 has shown is that there are PTs that are being passed from composition to composition. Whether this copying of PTs is a consequence of their longevity is unfortunately unclear.

7.4 Evidence for Fecundity

The fecundity of a meme relates to the frequency with which a meme instance is copied, and as a consequence, how abundant the meme becomes. For the purposes of the present research, fecundity relates to the propagation of PTs. For example, a PT that is being replicated on a yearly basis will have a far greater chance of becoming dominant in the pool of PTs than one that is only replicated once every fifty years: the former will have a greater fecundity than the latter. The more fecund a PT is, the more dominant within the pool of PTs it will become with respect to less fecund PTs, and therefore the former will have a greater possibility of selection, replication and variation over the latter. So fecundity, like longevity, can be related to how abundant a PT is.

As with longevity, it is not possible directly to investigate evidence for fecundity from the dataset. However, like longevity, it is possible to infer fecundity from the relative abundance of a PT. The ranking positions show that a PT which appears towards the top of the ranking positions within a particular period has a greater dominance, and therefore abundance, within that period than a PT that appears towards the bottom of the ranking positions. This can be translated as meaning that the dominant PTs have been selected and replicated more widely than those that appear at the bottom of the ranking positions, as argued in Section 6.2.1 above. As such, if a PT is moving upwards through the ranking positions across the periods, then it is displaying selection and replication, and consequently exhibiting higher fecundity than other PTs, because the former has been more widely copied than the latter in order to produce a higher ranking position.

Table 7.1 above shows certain PTs that continuously progress upwards through the ranking positions across all the periods. These PTs display greater abundance than those lower down the ranking positions, with the majority of the PTs in the table appearing in 50% or more of the compositions in the final period. However, there are some PTL/PTD combinations (the seven-note AD-PT, and the eleven-note A-PT and AD-PT) where no PT progresses continuously upwards through the ranking positions across all the periods. Additionally, not all of the PTs in this table appear in the majority of compositions in period 5 (for example, the seven-note RD-PT only appears in two out of the eight compositions in period 5).

Table 7.1, therefore, presents rather inconclusive evidence for fecundity. Some of the PTs do show a greater abundance in period 5 than in period 1. This greater abundance means that they are being selected and replicated more widely than those PTs further down the ranking positions and, by extension, are displaying greater fecundity. However, the fact that there are PTL/PTD combinations that have no associated PTs, as well as the fact that not all of the PTs in the table show a significantly greater abundance in period 5 than in period 1, calls into question whether fecundity has been demonstrated.

The same arguments used in Section 7.3 concerning a composer's awareness of PTs, also apply to fecundity. The question of how much intentionality there is on the part of the composer when creating PTs, and of how much of an impact the methodology for generating PTs has had on the ranking positions, cannot be addressed by looking solely at these ranking positions. As with longevity, this means it is difficult to determine whether or not the abundance of certain PTs is a consequence of fecundity.

7.5 Evidence for Copying-Fidelity

Perhaps the most difficult of the replicator properties to address is copying-fidelity. This attribute of a meme relates to the accuracy of the replication process. If a meme is replicated accurately every time, then it exhibits a high degree of copying-fidelity. Conversely, if a meme is copied with a degree of variation every time, then it exhibits low copying-fidelity. Consequently, if a meme is replicated regularly with a variation every time, then the abundance of the original meme will not be as great as that of a meme that replicates with the same frequency but with only occasional, or indeed, no copying errors. As such, it should be possible to infer whether a PT has shown copying-fidelity by its abundance across the periods, as long as some are shown to have descendants. However, the difficulty lies in determining whether a PT has any descendants (see Section 6.3 above).

Using the ranking tables, it is possible to determine whether a particular PT has greater abundance than other PTs. It is not possible, however, to use the ranking tables to determine if an abundant PT has high or low copying-fidelity. This is because the ranking tables do not show any form of ancestral relationship between PTs across periods. Consequently, it is not possible to prove that a PT with greater abundance than other PTs has a high level of copying-fidelity.

Nevertheless, by looking at some individual PTs in terms of their possible descendants (from the similarity algorithm) and determining whether those descendants and the original PTs are successful in the later periods (by looking at the ranking positions) it may be possible to infer that high copying-fidelity has been exhibited. If the original PT is still successful in the later periods then it can be argued that it is exhibiting high copying fidelity, in that while it has been varied at some stage in its life span (because there is a potential descendant) the original form also remains abundant.

Because copying-fidelity relates to the production of variants, it could be expected that the most abundant PTs will have produced fewer variants than the less abundant ones in order to remain the most abundant PTs. The highest ranking positions for each period show which PTs are the most abundant and, conversely, the bottom ranking positions show which are the least abundant. The highest ranking PTs, because they are abundant, should show fewer possible variant PTs than their bottom position counterparts if high copying-fidelity is to be inferred. However, there is the problem of taking into account the relative frequencies of the PTs. For example, the highest ranking position PTs will have many more occurrences than the lowest ranking position PTs and subsequently the former will have a greater opportunity to create variants than the latter.

Table 7.3 (on p. 235) shows the top (T) and bottom (B) ranking position PTs across all the periods together with the number of possible variant PTs (according to the similarity algorithm) they produce which do not appear in period 1.

A mixed picture is shown in Table 7.3 in regard to providing evidence for copying-fidelity. For example, the majority of the PTs at the top of the ranking positions have more possible variants than those at the bottom of the ranking positions. When looking at the difference between the number of possible variants between the top and bottom ranking position PTs, neither the PTD nor the PTL seems to affect whether the top or the bottom PT has the greater number of possible variants.

Whilst Table 7.3 shows that there are PTs that are more abundant than others, the correlation between a PT’s abundance and how many possible variants it spawns seems at odds with the logic that the more abundant PTs should have fewer antecedents than the less abundant PTs. This makes it difficult to prove whether high copying-fidelity is evident.

The drawbacks regarding trying to infer longevity and fecundity from the ranking tables also apply to copying-fidelity. It is possible that a composer will alight upon a particular PT completely independently of other composers. This would mean that the ranking tables could not be used to infer abundance, because the assumption is that the PTs are being copied between composers. There is also the problem that the methodology can generate unintentional PTs by combining musical phrases together, generating PTs that a composer would not consider as meaningful patterns.

Like longevity and fecundity, copying-fidelity is therefore difficult to prove. It has to be inferred from the ranking tables through the idea that if a PT displays abundance then it must have been copied repeatedly with few or no errors in the copying process. Table 7.3 provides a mixed in this respect.

It is also difficult to prove conclusively that two PTs are variants of each other within the dataset, as shown in Section 6.3 above. Consequently, it is doubtful whether high copying-fidelity has been, or indeed can be, proved using the current methodology.

PTL	PTD	T or B	PT	Ranking Position in Period					Nº of Possible Variants
				1	2	3	4	5	
3	A-PT	T	G/D0F/D0E	5	1	1	11	6	270
		B	E-1/D1E-1/U0E	747	1,318	1,222	1,003	511	81
	AD-PT	T	G:1/D0F:1/D0E:1	14	4	1	10	8	220
		B	C:1/U0F:2/D0C:1	593	1,214	1,008	778	394	22
	R-PT	T	0/D2/D1	1	1	1	1	4	8
		B	0/D12/U21	392	501	538	488	367	14
	RD-PT	T	0:1/D2:1/D1:1	1	1	1	1	5	8
		B	0:1/D4:1.3333/U4:1	615	1,086	954	769	439	4
7	A-PT	T	D/S0D/S0D/S0D/S0D/S0D /S0D	4	2	3	2	1	6,561
		B	B/00R/U0C/00R/U0D/00R /U0E	193	493	295	141	111	147
	AD-PT	T	D:1/S0D:1/S0D:1/S0D:1 /S0D:1/S0D:1/S0D:1	4	2	6	2	1	3,140
		B	E:1/S0E:1/S0E:1/S0E:1 /S0E:1/S0E:2/S0E:1	138	378	198	102	78	107
	R-PT	T	0/S0/S0/S0/S0/S0/S0	1	1	1	1	1	767
		B	0/R0/U1/R0/D1/R0/U7	372	780	632	396	234	794
	RD-PT	T	0:1/S0:1/S0:1/S0:1/S0:1 /S0:1/S0:1	1	1	1	1	1	449
		B	0:1/S0:1/S0:1/S0:1/S0:1 /S0:1/U12:2	241	556	472	262	162	7
11	A-PT	T	D/S0D/S0D/S0D/S0D/S0D /S0D/S0D/S0D/S0D/S0D	3	3	3	1	2	6,561
		B	G/00R/U0C/00R/S0C/00R /S0C/00R/S0C/00R/S0C	75	163	91	70	44	186
	AD-PT	T	D:1/S0D:1/S0D:1/S0D:1 /S0D:1/S0D:1/S0D:1/S0D:1 /S0D:1/S0D:1/S0D:1	3	2	6	2	2	3,140
		B	E-1:1/00R:1/S0E-1:1 /00R:1/S0E-1:1/00R:1 /S0E-1:1/00R:1/S0E-1:1 /00R:1/S0E-1:1	66	136	71	46	30	3,140
	R-PT	T	0/S0/S0/S0/S0/S0/S0/S0 /S0/S0/S0	1	1	1	1	1	767
		B	0/S0/S0/U1/S0/S0/S0/S0 /S0/S0/U1	148	323	255	175	93	40
	RD-PT	T	0:1/S0:1/S0:1/S0:1/S0:1 /S0:1/S0:1/S0:1/S0:1/S0:1 /S0:1	1	1	1	1	1	449
		B	0:1/R0:1/S0:1/R0:1/S0:1 /R0:1/D3:1/R0:1/S0:1/R0:1 /S0:1	123	261	210	118	62	92

Table 7.3: Top and bottom ranking position PTs with the number of possible variants that do not appear in period 1

7.6 Summary

It is difficult to provide evidence for the replicator properties of longevity, fecundity and copying-fidelity owing to these properties relating to the individual instances of the memes rather than the meme as a global entity, and the data coming from one manifestation of one storage mechanism. Therefore, a different approach was taken in that it was argued that these properties could be inferred from the abundance of the memes within the sample pool. This is because in order for some memes to become more abundant than others, they needed to manifest greater longevity and fecundity, and higher copying-fidelity than those that are less abundant.

There is evidence within the boundaries of the present study for some PTs becoming more abundant across the periods in terms of the increasing number of PTIs and the number of movements in which they appear. However, not all of the PTD/PTL combinations show such an increase in abundance across the periods when taking into account the total number of compositions in each period. Therefore, the evidence for the greater abundance of some PTs over others is inconclusive and, consequently, the evidence for the replicator properties of longevity and fecundity is also inconclusive.

Inferring evidence from the abundance of PTs for high copying-fidelity was not as straight-forward as for longevity and fecundity. This was due to copying-fidelity requiring the variation of PTs to be measured, i.e., if a PT shows a greater abundance than others, then it also needs to be shown that the PT has spawned fewer variants. Therefore, the difficulty in determining when PTs are variants of each other also affects the search for evidence of high copying-fidelity.

To provide evidence for high copying-fidelity, the most and least abundant PTs across all periods were investigated for the number of potential variants using the similarity algorithm. The argument was that the most abundant PTs should show fewer potential variants than the least abundant ones, because the former should have produced fewer potential variants in order gain their greater abundance. Unfortunately the data provided a mixed picture, with some PTs providing evidence for the hypothesis and others suggesting the contrary.

Part IV - Conclusions

8 Chapter 8: Do Memes Exist in Music?

8.1 Introduction

The ultimate aim of the present research was to provide evidence to support the hypothesis that memes exist in music. In order for evidence to be garnered, investigations into what constitutes a meme in music were undertaken, which led to a working definition of a meme in music. From this definition, a methodology was devised that allowed a computer program to search for evidence of memes using mass data analysis. This final chapter provides a review of the processes involved in finding evidence for memes in music, the results of the data analysis, and suggestions for further improvements to both the data and methodology to support the case for memes existing in music.

8.2 Defining the Meme

Unfortunately there is no clear definition of what constitutes a meme. Chapter 1, Section 1.3 above showed that a number of commentators support different definitions of a meme (for example Aunger (2000), Borenstein (2004), Brodie (1996), Gabora (1997), etc.) that concentrate on different aspects of the meme, such as the storage mechanism or the replication process. Additionally, no meme has been identified to the general acceptance of the academic community, a point made by Aunger (2002, p. 21). Therefore, a generic definition based on Gabora's definition of a meme (1997)was used in the present research: *a meme is a unit of cultural information that evolves by means of natural selection.*

Section 2.2 above showed that there are three main processes involved in natural selection: selection, replication and variation. Additionally, Dawkins argues that there are three replicator properties that are essential for memes to exist: longevity, fecundity and copying-fidelity (1989, pp. 193-194). These evolutionary processes and replicator properties can therefore be used to help verify the existence of memes. If a unit of cultural information exhibits these processes and properties, then that unit could arguably be considered a meme.

When applying the meme concept to music the main problem is identifying *a unit of cultural information* in music. A pattern in music would seem to be the obvious choice because it holds musical information. However, trying to determine what constitutes a pattern in music is problematic. For example, a pattern can be structural (such as a formal schema), polyphonic, harmonic, rhythmic, melodic, etc. There are additional problems with defining patterns in terms of identifying their boundaries and importance within compositions, as well as issues surrounding the involvement of secondary parameters such as metre and dynamics, and devices such as ornamentation. All of these issues make using patterns as the *unit of cultural information* problematic.

Despite these problems, a definition of a pattern in music was arrived at, against which the evolutionary processes and replicator properties could be tested. The definition was based on Miller's assertion that the brain chunks information into 'seven plus or minus two' units of information (1956), together with Jan's hypothesis that a meme cannot be shorter than three notes, in order to provide it with enough information to be distinct and salient (2007, pp. 60-61). Therefore the pattern in music, for the purposes of the present research, was defined as *any three to eleven monophonic consecutive notes, excluding symbolic ornamentation and secondary parameters.*

Consequently, the working definition of a meme in music used in the present study was:

Any three to eleven monophonic consecutive notes, excluding symbolic ornamentation and secondary parameters, that evolves by means of natural selection (i.e., selection, replication and variation), and which exhibits the replicator properties of longevity, fecundity and copying-fidelity.

8.3 The Methodology

The next stage of the research was to devise a methodology that could supply evidence for memes in music based on the working definition. This methodology needed to be able to show, over a period of time, the evolutionary processes taking place in music, as well as evidence for the replicator properties. In order to provide this evidence, a mass data analysis of music scores was undertaken using Knowledge Discovery in Database (KDD) techniques on a relational database model. The relational database model was selected to perform the data analysis because the technology is designed specifically for data storage and manipulation, has in-built indexing algorithms, and uses the Structured Query Language (SQL).

Elmasri and Navathe's six stages for KDD (2007, pp. 946-947) were used as a framework for the mass data analysis. The first stage was data selection, which involved sourcing scores in MusicXML format and importing these documents into the relational database. It was decided to use string quartets owing to the number of scores available in MusicXML format, and the ease with which it would be possible to generate patterns according to the working definition. Unfortunately, the number and range of scores available in MusicXML (or Kern) was not sufficient to provide meaningful results (i.e., it did not cover enough composers across a suitable time-span). Therefore additional scores were converted to MusicXML using *Photoscore* and *Sibelius* using both pdfs available online and scanned printed scores. In total, 442 movements from 111 compositions by 19 composers were used, ranging from early Haydn to late Shostakovich.

The next stage of KDD was data cleansing, which involved correcting errors in the data. There were a number of errors found in the documents sourced for the research. There were problems with the conversion from printed scores, editorial errors, the conversion from Kern to MusicXML, andusing *Photoscore* to convert to MusicXML. Where any errors were found, they were corrected in either *Photoscore* or *Sibelius*. However, owing to the intensive nature of the task, it could not be guaranteed that all errors were found and corrected.

Following on from data cleansing was the process of data enrichment, which involved adding extra data to facilitate the analysis. Two extra pieces of information were added to the database to this end. Firstly, a table was added that linked all enharmonically equivalent pitches together. This was necessary to expedite the calculations regarding intervallic distances between notes. Secondly, information was added in order to group the compositions together within predefined periods. This involved entering information on when a composition was begun, as well as information on which compositions should be grouped together. In the end, five periods were created, enabling the system to trace patterns across different time frames.

The next stage involved data transformation, where the data was manipulated using stored procedures to enhance the process of data mining. This involved calculating the periods; flagging the start and end of pieces, movements and instrumental parts; deciding which notes to use from chordal passages; converting tied notes of the same pitch to a single note; combining consecutive rests into one rest; working out the absolute value, relative intervallic value and relative durational value of each note compared to the previous note; generating all the possible three- to eleven-note patterns; and, finally, determining the pitch and duration properties of all the patterns generated.

After the data transformation had taken place, the data mining stage was implemented. This involved determining the frequency of occurrence of each pattern generated within each time period, according to the number of instances of the pattern and the number of movements in which it appeared, and then ranking the patterns according to this calculated frequency. Four sets of ranking positions were produced: absolute pitch values without duration (A-PT), absolute pitch values with relative durations (AD-PT), relative intervallic values without duration (R-PT), and relative intervallic values with relative durations (RD-PT).

Finally, the data reporting stage was implemented. This involved writing further stored procedures to interrogate the ranking positions and pattern properties, as well as using the similarity algorithm to help determine potential variation between patterns. The results were used to provide evidence of the evolutionary processes, as well as the replicator properties.

8.4 The Evidence

The ranking positions were investigated because tracking the position of pattern types across the periods was hypothesised to provide evidence for selection and replication. This was based on the argument that pattern types at the top of the ranking positions are more numerous than those at the bottom. Therefore, those at the top must have been selected and replicated more frequently than those at the bottom. Additionally, if certain pattern types showed a continuous progression upwards, or downwards, across the periods, then they were either being selected and replicated more widely, or were gradually declining in the frequency of their selection and replication, according to their progression direction. There were pattern types in the data that continuously progressed upwards or downwards across the periods for all pattern type lengths and descriptors. Additionally, there were pattern types that did not appear in the first, or last, period that continuously progressed upwards, or downwards, over the remaining periods. It was also shown that the range of pattern types appearing in a particular period was different for each of the periods, with over 90% of pattern types appearing in one period only, for all pattern type descriptors. However, it was shown that the longer pattern types had relatively more pattern types appearing in only one period than was the case with their shorter counterparts.

The investigation of variation was more problematic because the ranking positions by themselves could not provide suitable direct evidence. As a result, a similarity algorithm that compared the pitch and/or duration properties of the pattern types was used to determine if two pattern types could potentially be related. These relationships were investigated by looking at the pattern types that first appeared in period 3 and determining if any pattern types in the previous two periods could have been antecedents. If there were possible antecedents then this was interpreted as supporting evidence for variation. Additionally, the new pattern types in period 3 were then investigated to explore whether they showed signs of selection and replication in the final two periods. If it could be shown that some of the new pattern types in period 3 had possible antecedents in periods 1 and 2 (i.e., that they exhibited variation), as well as undergoing selection and replication in periods 4 and 5, then those pattern types could possibly be memes.

It was found that the number and range of pattern types in each period was different. Additionally, there were more new pattern types in each of periods 2 to 5 than pattern types that had appeared in any of the previous periods respectively. Both of these points were used to argue that there are pattern types new in periods 2 to 5 that could potentially be variants of extant pattern types. When investigating the pattern types in period 3 that had not appeared in either period 1 or 2, it was shown that the similarity algorithm produced a number of possible ancestral connections between the chosen period 3 pattern types and pattern types in periods 1 and 2. It was then shown that some of these new pattern types from period 3 progressed upwards through the ranking positions in periods 4 and 5. Therefore, the data arguably showed evidence within the boundaries of the present study for selection, replication and variation: there were some new pattern types in period 3 with possible antecedents that then progressed upwards through the ranking positions.

An alternative approach was used to provide evidence for the replicator properties. This was because the dataset only relates to one manifestation of one storage mechanism for memes: the dataset only points to a meme being found in a particular edition of a composition without regard to any other of its other manifestations. Consequently, it was not possible to measure the different meme manifestations for differences in their replicator properties.

Instead, it was argued that the replicator properties could be inferred from the relative abundance of memes. This was based on Dawkins' argument that the same replicator properties help to determine the abundance of genes (1989, pp. 16-17). For example, a gene with high longevity will not only have more chances for replication, but will also have more instances than a gene with low longevity. A gene with a high fecundity will replicate more frequently than a gene with low fecundity, again increasing the number of instances. Finally, a gene with high copying-fidelity will create more exact copies of itself than one with low copying-fidelity, again increasing the number of its instances. Therefore, in order for a meme to become more abundant than its rivals, it must have exhibited greater longevity and fecundity, and higher copying-fidelity, than them. Consequently, in this study the ranking positions were used to show if any pattern types had a greater abundance over others within the dataset.

The ranking positions produced evidence that there were pattern types with a greater abundance than others in the dataset owing to the ranking positions being based on the number of appearances of pattern types in the dataset. However, it was found that certain pattern type descriptor and length combinations did not have any pattern types that continuously progressed upwards through the ranking positions. Additionally, some of the pattern types that had the highest upward progression through the ranking positions appeared in 50% or less of compositions in the last period. Both of these points called into question the strength of the evidence for the replicator properties because, in certain cases, the relative abundance of pattern types could not clearly be determined.

Unfortunately, looking at the abundance of pattern types by itself was not enough to provide evidence for high copying-fidelity. This was because evidence needed to be found that variation had actually taken place. Therefore, the number of potential variant pattern types was calculated using the similarity algorithm for the pattern types that were either the most, or least, abundant across all the periods. However, the data provided a mixed picture in that, on the whole, there were a larger number of potential variants for the most abundant than for the least abundant pattern types, contrary to what would perhaps be expected with high copying-fidelity.

A couple of issues were raised when looking at some examples of the pattern types within the scores. Firstly, most of the examples showed the pattern types as part of a longer phrase or motive, often starting on the weak beat or part of a beat. Secondly, there was no evidence for pattern types being used in similar scenarios when they were more prominent in the ranking positions, or for them being used in diverse scenarios when they were less prominent. Consequently, there was no evidence for the placement of the pattern types having a bearing on a pattern type's prominence in the ranking positions.

The data presented for both the evolutionary processes and replicator properties shows that there are differences between each of the four PTDs as well as the different PTLs when investigating the evidence for memes in music. For example, all four three-note PTDs have PTs that continuously progress upwards through the ranking positions, whereas the seven-note AD-PT, and the eleven-note A-PT and AD-PT have none. There are also differences between the number of periods in which a given PT has a PTI. For example, for the three-note PTs, the R-PT has the highest (34.5%), and the AD-PT has the lowest (3.3%) percentage of PTs appearing in all five periods compared to the total number of PTs for each PTL/PTD combination. However, for the eleven-note PTs, none of the four

PTDs has enough PTs appearing in all five periods to register a percentage at one decimal place (i.e., 0.0%). This difference in the percentages of PTs appearing in all five periods across the different PTD/PTL combinations is also reflected in the number of new PTs arriving in each period. For example, for three-note PTs in period 5, the AD-PT has the highest (48%) and the R-PT has the lowest (7%) percentage figure for the number of new PTs compared to the total number of PTs in each period for each PTL/PTD combination, whereas for the eleven-note PTs, all four PTDs have a percentage figure of 95% or above.

There were also differences between the different PTL/PTD combinations in terms of the similarity algorithm. For example, the A-PTs and AD-PTs produced a greater number of possible relationships than the R-PTs and the RD-PTs. This was to be expected, because the A-PTs and the AD-PTs encompass all possible transpositions, together with their enharmonic equivalents pitches of their PTs as possible relations. Additionally, when comparing the possible antecedents to the PTs new in period 3, the seven- and eleven-note PTs produced, on the whole, a greater number of improbable matches compared with the three-note PTs, across all four PTDs.

8.5 Ten Potential Ideas for Further Investigations

Although the present study has produced some interesting results, and shown that mass data analysis of music is possible using relational databases, further work in this area could be undertaken to provide further support for the concept of memes in music. Ten ideas are highlighted that would potentially advance the study of memes in music using computer-based analysis techniques.

8.5.1 Enhancing the Dataset

Although there were 442 movements from 111 compositions by 19 different composers used in the present research, there were some notable absences from the list (such as Boccherini and Schoenberg). Additionally, owing to the availability of scores in a suitable format, there was a heavy reliance on string quartets by Haydn, Mozart and Beethoven, and this skewed the relative quantity of data towards the earlier periods (although it can be argued that this period was the height of string quartet writing). Therefore, a more comprehensive set of data that includes more composers from the later periods, such as Villa-Lobos and Schoenberg, could be created. The ranking positions could then be recreated and re-analysed for more compelling evidence of the evolutionary processes and the replicator properties. However, adding additional string quartets will still mean that the dataset is using only a small subset from the whole music corpus.

8.5.2 Redefining the Pattern

In the present study, the pattern generation algorithm created all three- to eleven-note patterns regardless of their placement in the music. This inevitably produced the situation where patterns would be generated using the middle of phrases, or across natural phrase boundaries such as cadence points, etc. Therefore, it can be argued that the majority of patterns generated by the system were not meaningful from a musical perspective. Additionally, the pattern generation algorithm resulted in relatively small figures for patterns displaying possible evidence for memes compared to the total number of patterns generated by the dataset. An improvement would be to create an algorithm that looks for patterns from a more specifically musical perspective (e.g., aligning pattern generation with musical phrases, or using the Gestalt principles of proximity, similarity and good continuation as the basis for pattern generation), rather than all possible patterns. Again, the ranking positions could then be re-analysed to determine whether more compelling evidence for the evolutionary processes and replicator properties exists.

8.5.3 Investigating Differences in Pattern Type Lengths

It may also be possible to create an algorithm to relativize the different length pattern types. This would then enable them to be compared within a single ranking system. It has already been stated that the problem of comparing different pattern lengths is similar to the concept of relative species abundance (Section 5.5). Therefore, it may be possible to utilise some of the research into relative species abundance to create a suitable algorithm to compare patterns of different length. This would allow further investigation into the impact of the length of pattern types on the meme (i.e., would three-note patterns still dominate the top ranking positions).

Additionally, the analytical component of the present research concentrated on three-, seven- and eleven-note pattern types. This showed that there was a difference in the evidence obtained between the three categories. Clearly, the intervening pattern type lengths could also be investigated. This investigation would be able to show if a proportional relationship exists between the pattern type length and the nature of the evidence provided for memes.

8.5.4 Enhancing the Similarity Algorithm

As explained in Section 3.3 above, it is intrinsically difficult to determine the extent to which two patterns are similar to each other. Methods exist to understand connections between patterns but their implementation was beyond the scope of the present research. Researchers have used techniques such as weighting systems (Rolland, 1999) and neural networks (Sotiropoulos, et al., 2008) to help determine similarity between patterns, which could be beneficial in creating an enhanced similarity algorithm for the present research project. Additionally, there are other approaches, such as using web-search algorithms based on vectors (Haveliwala, 2003), which could also be investigated to determine whether they are suitable for music pattern matching. Appropriate techniques could then be mapped on to the present research to provide a more sophisticated method for determining whether patterns are ancestrally related, allowing for more compelling evidence for variation to be gathered.

8.5.5 Investigating Replicator Property Impact

The present research has shown that some pattern types are more abundant than others, but was unable to determine which of the three replicator properties had the greatest impact on this abundance. Further investigation would involve generating details on other manifestations of memes, such as other editions, performances and recordings. This will be a time-consuming task, owing to the diverse nature of these different manifestations. However, the details of these forms could be analysed (e.g., for a recording, using details as to the length of availability and the number distributed for each recording) to provide supporting evidence for the replicator properties as well as to determine which property is the most consequential.

8.5.6 Using Scale Degree-Pattern Generation

Rather than looking at patterns without regard to the tonal context of the surrounding music, as was the case with the present research, the system could be developed to use scale degree mapping as well. The advantage of using scale degrees would be that patterns with a similar tonal context could be compared. However, an algorithm to determine scale degrees from the pitches' tonal context would be required. This would be relatively straightforward for earlier, more traditional works, but more involved for later, more complex and diverse compositions. It would, nevertheless, be possible in principle to compare new and original results to determine if, and to what extent, tonal context is important to the concept of memes in music.

8.5.7 Redefining the Time Periods

Currently, the system divides the data into five different time periods. These could be redefined to investigate how much of an impact they have on the results, or to investigate the evolution of certain patterns over different time periods. It would even be possible to create time periods for an individual composer, to determine if there was a change in their use of patterns over time. For example, periods could be set up to match the traditional three stylistic periods of Beethoven's output.

8.5.8 Adding Geographical Detail

Another way to refine the data would be to take into account the geographical location of composers. In the same way that compositions are linked in the present study to a time period, it would be possible to associate them with geographical regions. It could then be investigated whether certain geographical regions showed greater evidence for the replication of memes than others. This would involve creating a larger dataset to ensure that there was enough data in each region and time period combination to make the results meaningful.

8.5.9 Investigating Composers' Lineages

It would also be possible to compare lineages of composers who are intrinsically linked. These lineages could then be compared to the dataset as a whole. This comparison could show whether there is more sharing of patterns within certain lines of composers than in the dataset as a whole.

8.5.10 Separating Out Instrumental Parts

In the present research, all the instrumental parts were grouped together when analysing the ranking positions. It would be possible to separate out the different instrumental parts, and then reanalyse the ranking positions. This would help determine if there is more evidence for memes, or if certain patterns are found more often, in a particular instrumental part over others. It would also help to determine whether there is a stylistic change in how instrumental parts are utilised within the string quartet over time.

8.6 Three Alternative Investigations

The previous section has highlighted ten possible areas of investigation that could help extend the present study. This section briefly examines three areas that have not been fully explored in the present study but which may be useful when investigating memes in music.

8.6.1 The Role of Consciousness

Although the issue of how memes are potentially stored in the brain was raised in Section 2.5.1, the issues surrounding the role of memes in the conscious were deliberately avoided. This was because the methodology used was designed only to determine if the evolutionary processes were taking place, and not how they were driven. Additionally, the whole subject of how the conscious works within music is extremely complex. Additionally, there is the problem that composers will consciously try and find their own voice through their music by creating their own styles and patterns, which is counter-intuitive to the idea that memes are passed between composers. Indeed, a whole thesis could be undertaken on the effect of consciousness on the meme in music.

8.6.2 Memetic Signatures

Another aspect that has not been fully explored in the present study is what constitutes the environment when investigating memes in music. Darwin argued that the environment plays a significant factor in evolution, in that it helps shape the survival of some traits over others (1859 repr. 1985, pp. 344-396). Therefore, are there factors, both within and surrounding music, which have a bearing on the selection and replication of memes?

Such factors would include the composer, the performer, listener, instruments, and the genre. Each of these could have their own set of restrictions (e.g., the composer by their imagination, the performer by their skill, the listener by their concentration, the instruments by their range, the genre by its instruments, etc.). Unfortunately, it is not possible to investigate all of these factors through the dataset. However, it is possible to use the current dataset to investigate the differences between the patterns used by different composers, or common on particular instruments. It might also be possible to devise an algorithm that creates a 'memetic signature' for each composer, and/or each instrument, and/or each string quartet, giving each one a unique identifier that can also be used as a measure of comparison. This signature could then be used in areas such as concert planning (i.e., which compositions sit well together), and teaching (i.e., which pieces have a natural progression of difficulty), etc.

8.6.3 Pattern Salience

Another issue that impacts on the consciousness debate is that of the salience of patterns. Again, this is an issue that has not been investigated owing to the complexity of providing a general definition of salience that can be used in the context of the present study. However, what has been shown is that the examples used in the present study, on the whole, tended to be part of longer phrases, and tended to commence on weak metric positions. This therefore calls into question whether the examples used exhibit sufficient salience to meet Jan's criteria for a meme (see p. 54).

The range of examples used in the present study was relatively small compared to the number of patterns generated by the system. Therefore, it is questionable how representative the examples were. Nevertheless, further investigations could be undertaken to determine if there is any connection between the relative salience of the patterns and their ranking positions. If a link is discovered, then this could add support to Jan's hypothesis that '[t]hose memes that have the greatest perceptual-cognitive salience will tend to be the most widely propagated in the meme pool, to the disadvantage of those less salient forms' (2007, p. 229). However, the drawbacks to such investigations are providing a suitable definition for salience, and sifting through, analysing and categorising large numbers of examples.

8.7 Summary

In order to investigate the concept of memes in music, it was naturally important to understand what a meme is. Dawkins hypothesised the meme/gene analogy as a way of explaining the mechanisms behind cultural evolution. Unfortunately, there is little agreement on what constitutes culture and consequently the concept of the meme is similarly loose. However, if it is accepted that the meme evolves by means of natural selection in an analogous way to that of genes, then finding evidence for selection, replication and variation in culture will help provide evidence that memes exist. Additionally, Dawkins draws on gene replication to argue that memes must also exhibit longevity, fecundity and copying fidelity; if these properties are exhibited then this constitutes further evidence for the existence of memes.

Using the ranking positions and the similarity algorithm, evidence was found for selection, replication and variation within the corpus of music used for the present research. However, the data for the progression of pattern types across the ranking positions was limited compared to the total number of pattern types in the dataset. Longevity, fecundity and copying-fidelity were inferred from the abundance of certain pattern types over others, using the ranking positions. However, when investigating the abundance of pattern types in relation to the total number of compositions in the dataset, the evidence became less convincing than that for the evolutionary processes. In conclusion, although evidence supporting the possible existence of memes in the corpus of music used for this study has been provided within the confines of the definitions and assumptions stipulated in this thesis, additional work is required to help advance the case for the existence of memes in music.

Appendices

Appendix 1 - Working Definitions

Meme: a unit of cultural information that evolves by means of natural selection.

Meme in Music: Any three to eleven monophonic consecutive notes, excluding symbolic ornamentation and secondary parameters, that evolves by means of natural selection (i.e., selection, replication and variation), and which exhibits the replicator properties of longevity, fecundity and copying-fidelity.

Pattern in Music: Any three to eleven monophonic consecutive notes, excluding symbolic ornamentation and secondary parameters.

Pattern Similarity: When patterns have an exact match on their calculated property values.

Appendix 2 - Repertoire List

Work	Mvts	Source	Downloaded	Encoder
Bartok, (n.d.) *String Quartet No. 1*. London: Boosey and Hawkes.	All	P	09/11/2010	AH
Bartok, (1939) *String Quartet No. 2*. London: Boosey and Hawkes.	All	P	09/11/2010	AH
Beethoven, (1801) *String Quartet Op. 18 No. 1*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1801) *String Quartet Op. 18 No. 2*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1801) *String Quartet Op. 18 No. 3*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1801) *String Quartet Op. 18 No. 4*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1801) *String Quartet Op. 18 No. 5.* Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1801) *String Quartet Op. 18 No. 6*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1808) *String Quartet Op. 59 No. 1*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1808) *String Quartet Op. 59 No. 2*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1808) *String Quartet Op. 59 No. 3.* Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1810) *String Quartet Op. 74*. Leipzig: Unknown.	All	K	08/11/2010	CS
Beethoven, (1816) *String Quartet Op. 95*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1826) *String Quartet Op. 127*. Mainz: Unknown.	All	K	08/11/2010	CS

Work	Mvts	Source	Downloaded	Encoder
Beethoven, (1827) *String Quartet Op. 130*. Vienna: Unknown.	All	K	08/11/2010	CS
Beethoven, (1827) *String Quartet Op. 131*. Mainz: Unknown.	All	K	08/11/2010	CS
Beethoven, (n.d.) *String Quartet Op. 132*. n.p.: Unknown.	All	K	08/11/2010	CS
Beethoven, (1970) *String Quartet Op. 133*. New York: Dover Publications.	All	P	09/11/2010	AH
Beethoven, (1827) *String Quartet Op. 135*. Berlin: Unknown.	All	K	08/11/2010	CS
Borodin, (1994) *String Quartet No. 1*. New York: Dover Publications.	All	P	29/11/2010	AH
Borodin, (1994) *String Quartet No. 2*. New York: Dover Publications.	All	P	29/11/2010	AH
Brahms, (1926) *String Quartet No.1*. Leipzig: Breitkopf & Härtel.	1	P	29/11/2010	AH
Brahms, (n.d.) *String Quartet No. 1*. n.p.: Unknown.	2,3	K	8/11/2010	CS
Brahms, (n.d.) *String Quartet No. 1*. n.p.: Unknown.	4	G	9/11/2010	GP
Brahms, (1926) *String Quartet No. 2*. Leipzig: Breitkopf & Härtel.	All	P	29/11/2010	AH
Brahms, (1926) *String Quartet No. 3.* Leipzig: Breitkopf & Härtel.	All	P	29/11/2010	AH
Debussy, (1987) *String Quartet Op. 10*. Mineola: Dover Publications.	1	P	29/11/2010	AH
Debussy, (n.d.) *String Quartet Op. 10*. London: Eulenberg.	2,3,4	S		AH
Dvořák, (n.d.) *String Quartet Op. 96*. London: Eulenberg.	All	S		AH
Dvořák, (n.d.) *String Quartet Op. 105*. London: Eulenberg.	All	S		AH
Grieg, (1991) *String Quartet Op. 27*. Mineola: Dover Publications.	All	P	10/05/2011	AH

Work	Mvts	Source	Downloaded	Encoder
Haydn, (1772) *String Quartet Op. 1 No. 1*. n.p.: Bremner Edition.	All	K	08/11/2010	FB
Haydn, (1772) *String Quartet Op. 1 No. 2*. n.p.: Bremner Edition.	All	K	08/11/2010	FB
Haydn, (1772) *String Quartet Op. 1 No. 4*. n.p.: Bremner Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 9 No. 3*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 17 No. 1*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 17 No. 3*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 17 No. 5*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 20 No. 1*. n.p.: Trautwein Edition.	1,2,4	K	08/11/2010	FB
Haydn, (1985) *String Quartet Op. 20 No. 1.* New York: Dover Publications.	3	P	07/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 20 No. 2*. n.p.: Trautwein Edition.	1,3,4	K	08/11/2010	FB
Haydn, (1985) *String Quartet Op. 20 No. 2*. New York: Dover Publications.	2	P	07/12/2010	AH
Haydn, (1985) *String Quartet Op. 20 No. 3*. New York: Dover Publications.	1	P	07/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 20 No. 3*. n.p.: Pleyel Edition.	2,3,4	K	08/11/2010	FB
Haydn, (1985) *String Quartet Op. 20 No. 4*. New York: Dover Publications.	1,2	P	07/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 20 No. 4*. n.p.: Trautwein Edition.	3,4	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 20 No. 5*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 33 No. 1*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB

Work	Mvts	Source	Downloaded	Encoder
Haydn, (n.d.) *String Quartet Op. 33 No. 3*. n.p.: Richault Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 33 No. 5*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 33 No. 6*. n.p.: Richault Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 42. n.p.:* Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 50 No. 2*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 50 No. 3*. n.p.: Trautwein Edition.	1,3,4	K	08/11/2010	FB
Haydn, (1982) *String Quartet Op. 50 No. 3*. New York: Dover Publications.	2	P	07/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 50 No. 4*. n.p.: Trautwein Edition.	1,3,4	K	08/11/2010	FB
Haydn, (1982) *String Quartet Op. 50* No. 4. New York: Dover Publications.	2	P	07/12/2010	AH
Haydn, (1982) *String Quartet Op. 50 No. 5*. New York: Dover Publications.	1,2	P	07/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 50 No.* 5. n.p.: Trautwein Edition.	3,4	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 54 No. 1*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 54 No. 2*. n.p.: Dover Publication.	1,3,4	K	08/11/2010	FB
Haydn, (1982) *String Quartet Op. 54 No. 2*. New York: Dover Publications.	2	P	07/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 54 No. 3*. n.p.: Trautwein Edition.	All	K	08/11/2010	FB
Haydn, (1980) *String Quartet Op. 55 No. 1*. New York: Dover Publications.	1	P	07/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 55 No. 1*. n.p.: Dover Publication.	2,3,4	K	08/11/2010	WH

Work	Mvts	Source	Downloaded	Encoder
Haydn, (n.d.) *String Quartet Op. 55 No. 2*. n.p.: Dover Publication.	All	K	08/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 55 No. 3*. n.p.: Dover Publication.	1,3,4	K	08/11/2010	WH
Haydn, (1980) *String Quartet Op. 55 No. 3*. New York: Dover Publications.	2	P	07/12/2010	AH
Haydn, (1980) *String Quartet Op. 64 No. 2*. New York: Dover Publications.	1,2	P	08/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 64 No. 2*. n.p.: Dover Publication.	3,4	K	09/11/2010	WH
Haydn, (n.d.) *String Quartet Op. 64 No. 3*. n.p.: Dover Publication.	1,3,4	K	09/11/2010	FB
Haydn, (1980) *String Quartet Op. 64 No. 3*. New York: Dover Publications.	2	P	08/12/2010	AH
Haydn, (1980) *String Quartet Op. 64 No. 5*. New York: Dover Publications.	1	P	08/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 64 No. 5*. n.p.: Dover Publication.	2,3,4	K	09/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 71 No. 3*. n.p.: Trautwein Edition.	All	K	09/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 74 No. 1*. n.p.: Trautwein Edition.	All	K	09/11/2010	FB
Haydn, (1979) *String Quartet Op. 74 No. 3*. New York: Dover Publications.	1	P	08/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 74 No. 3*. n.p.: Trautwein Edition.	2,3,4	K	09/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 76 No. 1*. n.p.: Trautwein Edition.	1,2,3	K	09/11/2010	FB
Haydn, (1979) *String Quartet Op. 76 No. 1*. New York: Dover Publications.	4	P	08/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 76 No. 2*. n.p.: Trautwein Edition.	1,2,3	K	09/11/2010	FB
Haydn, (1979) *String Quartet Op. 76 No. 2*. New York: Dover Publications.	4	P	08/12/2010	AH

Work	Mvts	Source	Downloaded	Encoder
Haydn, (1979) *String Quartet Op. 76 No. 3*. New York: Dover Publications.	1,4	P	08/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 76 No. 3.* n.p.: Trautwein Edition.	2,3	K	09/11/2010	FB
Haydn, (n.d*.) String Quartet Op. 76 No. 4*. n.p.: Trautwein Edition.	All	K	09/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 76 No. 5*. n.p.: Trautwein Edition.	All	K	09/11/2010	FB
Haydn, (n.d.) *String Quartet Op. 76 No. 6.* n.p.: Trautwein Edition.	All	K	09/11/2010	FB
Haydn, (1979) *String Quartet Op. 77 No. 1*. New York: Dover Publications.	1	P	08/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 77 No. 1*. n.p.: Trautwein Edition.	2,3,4	K	09/11/2010	FB
Haydn, (1979) *String Quartet Op. 77 No. 2*. New York: Dover Publications.	1,3	P	08/12/2010	AH
Haydn, (n.d.) *String Quartet Op. 77 No. 2*. n.p.: Trautwein Edition.	2,4	K	09/11/2010	FB
Janáček, (1990) *String Quartet No. 1*. n.p.: Universal Edition	All	S		AH
Mendelssohn, (n.d.) *String Quartet Op. 12*. Mineola: Dover Publications.	All	P	06/12/2010	AH
Mendelssohn, (n.d.) *String Quartet Op. 44 No. 2*. Mineola: Dover Publications.	All	P	06/12/2010	AH
Mendelssohn, (n.d.) *String Quartet Op. 80*. Mineola: Dover Publications.	All	P	06/12/2010	AH
Mozart, (n.d.) *String Quartet K. 80.* n.p.: Breitkopf & Härtel	All	K	09/11/2010	EC
Mozart, (n.d.) *String Quartet K. 155*. n.p.: Breitkopf & Härtel	All	K	09/11/2010	EC
Mozart, (n.d.) *String Quartet K. 156.* n.p.: Breitkopf & Härtel	All	K	09/11/2010	EC
Mozart, (n.d.) *String Quartet K. 157*. n.p.: Breitkopf & Härtel	All	K	09/11/2010	EC

Work	Mvts	Source	Downloaded	Encoder
Mozart, (n.d.) *String Quartet K. 157*. n.p.: Breitkopf & Härtel	All	K	09/11/2010	EC
Mozart, (n.d.) *String Quartet K. 158*. n.p.: Breitkopf & Härtel	1,3	K	09/11/2010	EC
Mozart, (1881) *String Quartet K. 158*. Leipzig: Breitkopf & Härtel	2	P	07/12/2010	AH
Mozart, (n.d.) *String Quartet K. 159*. n.p.: Breitkopf & Härtel	1,2	K	09/11/2010	EC
Mozart, (1881) *String Quartet K. 159*. Leipzig: Breitkopf & Härtel	3	P	07/12/2010	AH
Mozart, (1881) *String Quartet K. 160*. Leipzig: Breitkopf & Härtel	1	P	07/12/2010	AH
Mozart, (n.d.) *String Quartet K. 160.* n.p.: Breitkopf & Härtel	2,3	K	09/11/2010	EC
Mozart, (n.d.) *String Quartet K. 168.* n.p.: Breitkopf & Härtel	All	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 169.* n.p.: Breitkopf & Härtel	1,3,4	K	09/11/2010	FB
Mozart, (1881) *String Quartet K. 169*. Leipzig: Breitkopf & Härtel	2	P	07/12/2010	AH
Mozart, (1882) *String Quartet K. 170.* Leipzig: Breitkopf & Härtel	1	P	07/12/2010	AH
Mozart, (n.d.) *String Quartet K. 170.* n.p.: Breitkopf & Härtel	2,3,4	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 171*. n.p.: Breitkopf & Härtel	All	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 172*. n.p.: Breitkopf & Härtel	1,2,4	K	09/11/2010	FB
Mozart, (1882) *String Quartet K. 172*. Leipzig: Breitkopf & Härtel	3	P	07/12/2010	AH
Mozart, (n.d.) *String Quartet K. 173*. n.p.: Breitkopf & Härtel	1,3,4	K	09/11/2010	FB
Mozart, (1882) *String Quartet K. 173*. Leipzig: Breitkopf & Härtel	2	P	07/12/2010	AH

Work	Mvts	Source	Downloaded	Encoder
Mozart, (n.d.) *String Quartet K. 387*. n.p.: Breitkopf & Härtel	All	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 421*. n.p.: Breitkopf & Härtel	All	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 428*. n.p.: Breitkopf & Härtel	1,2,3	K	09/11/2010	FB
Mozart, (1882) *String Quartet K. 428*. Leipzig: Breitkopf & Härtel	4	P	07/12/2010	AH
Mozart, (n.d.) *String Quartet K. 458*. n.p.: Breitkopf & Härtel	1,2,3	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 458*. n.p.: Unknown.	4	G	09/11/2010	GP
Mozart, (1882) *String Quartet K. 464*. Leipzig: Breitkopf & Härtel	1	P	07/12/2010	AH
Mozart, (n.d.) *String Quartet K. 464*. n.p.: Breitkopf & Härtel	2,3,4	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 465*. n.p.: Breitkopf & Härtel	All	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 499*. n.p.: Breitkopf & Härtel	1,3,4	K	09/11/2010	FB
Mozart, (1882) *String Quartet K. 499*. Leipzig: Breitkopf & Härtel	2	P	07/12/2010	AH
Mozart, (1882) *String Quartet K. 575*. Leipzig: Breitkopf & Härtel	1,4	P	07/12/2010	AH
Mozart, (n.d.) *String Quartet K. 575*. n.p.: Breitkopf & Härtel	2,3	K	09/11/2010	FB
Mozart, (n.d.) *String Quartet K. 589*. n.p.: Breitkopf & Härtel	1,3,4	K	09/11/2010	FB
Mozart, (1882) *String Quartet K. 589*. Leipzig: Breitkopf & Härtel	2	P	07/12/2010	AH
Mozart, (n.d.) *String Quartet K. 590*. n.p.: Breitkopf & Härtel	All	K	09/11/2010	FB
Prokofiev, (1948) *String Quartet No. 1*. New York: International Music Company	All	S		AH

Work	Mvts	Source	Downloaded	Encoder
Ravel, (1987) *String Quartet in F*. Mineola:Dover Publications	All	P	29/11/2010	AH
Schubert, (1965) *String Quartet No. 1*. New York: Dover Publications	All	P	06/12/2010	AH
Schubert, (1965) *String Quartet No. 7*. New York: Dover Publications	All	P	07/12/2010	AH
Schubert, (n.d.) *String Quartet No. 10*.	All	K	09/11/2010	CS
Schubert, (1965) *String Quartet No. 14*. New York: Dover Publications	All	P	07/12/2010	AH
Schumann, (n.d.) *String Quartet No. 1.*	All	G	09/11/2010	GP
Schumann, (1881) *String Quartet No. 2.* Leipzig: Breitkopf & Härtel	All	P	07/12/2010	AH
Schumann, (1881) *String Quartet No. 3.* Leipzig: Breitkopf & Härtel	All	P	07/12/2010	AH
Shostakovich, (1968) *String Quartet No. 3.* London: Eulenburg	All	S		AH
Shostakovich, (1961) *String Quartet No. 8.* London: Boosey and Hawkes	All	S		AH
Shostakovich, (1971) *String Quartet No. 13.* London: Boosey and Hawkes	All	S		AH
Sibelius, (1909) *String Quartet Op. 56*. London: Eulenberg	All	P	07/12/2010	AH
Smetana, (1946) *String Quartet No.1*. Prague: Hudební Matice	1,2	P	07/12/2010	AH
Smetana, (n.d.) *String Quartet No.1*. London: Eulenburg	3,4	S		AH
Tchaikovsky, (n.d.) *String Quartet No. 1*. New York: Dover Publications	All	P	07/12/2010	AH
Tchaikovsky, (1994) *String Quartet No. 3*. Mineola: Dover Publications	All	P	07/12/2010	AH

Key:

- Source
  - K – Downloaded from the Kern Score Website (http://kern.ccarh.org/)
  - P – Downloaded from the Petrucci Music Website (http://imslp.org/wiki/Main_Page)
  - G – Downloaded from the Project Gutenburg Website (http://www.gutenberg.org/wiki/Main_Page)
  - S – Scanned from a printed score
- Encoder
  - AH – Andrew Hawkett
  - CS – Craig Sapp
  - EC – Edmund Correia
  - FB – Frances Bennion
  - GP – Geoff Pawlicki
  - WH – Walter Hewlett

Appendix 3 - Stored Procedures

Stored Procedure Name	Action
Pass00:1Periods	Creates the periods in tblPiece
Pass01:1Parts	Flags the change from pieces, movements and parts
Pass02:1Chords	Flags the top notes in the violin 1 part chords
Pass02:2Chords	Flags the top notes in the violin 2 part chords
Pass02:3Chords	Flags the top notes in the viola part chords
Pass02:4Chords	Flags the bottom notes in the cello part chords
Pass02:5Chords	Flags all the notes that are not part of a chord
Pass02:6Chords	Flags all the unrequired notes of a chord
Pass03:1Ties	Combines tied notes and consecutive rests
Pass04:1NoteValues	Creates the relative movement between notes
Pass05:1Patterns	Generates all the patterns
Pass05:2Patterns	Removes unwanted patterns
Pass05:3Patterns	Recreates the indexes in tblPatterns
Pass06:0Index	Removes indexes from tblPatternProperty
Pass06:1:1Time	Creates time property records in tblPatternProperty
Pass06:1:2Time	Updates the TimeID column in tblPatterns
Pass06:1:3Time	Updates the Shape and PitchCentre columns in tblPatternProperty
Pass06:1:4Time	Updates the HighLow column in tblPatternProperty
Pass06:2:1Rel	Creates relative interval pattern property records in tblPatternProperty
Pass06:2:2Rel	Updates the RelID column in tblPatterns
Pass06:2:3Rel	Updates the Shape and PitchCentre columns in tblPatternProperty
Pass06:2:4Rel	Updates the HighLow column in tblPatternProperty
Pass06:3:1Abs	Creates absolute pitch pattern property records in tblPatternProperty
Pass06:3:2Abs	Updates the AbsID column in tblPatterns
Pass06:3:3Abs	Updates the Shape and PitchCentre columns in tblPatternProperty
Pass06:3:4Abs	Updates the HighLow column in tblPatternProperty
Pass06:4Index	Recreates the indexes on tblPatternProperty
Pass07:1Ranking	Creates ranking positions across all pattern type lengths
Pass07:2Ranking	Creates ranking positions within pattern type lengths
Pass07:3Ranking	Creates the pattern type and shape properties
Pass08:1DeDupA	De-duplicates A pattern descriptor pattern types
Pass08:1DeDupR	De-duplicates R pattern descriptor pattern types
Pass08:1DeDupAD	De-duplicates AD pattern descriptor pattern types
Pass08:1DeDupRD	De-duplicates RD pattern descriptor pattern types
Pass09:1Ranking	Creates new ranking positions across all pattern type lengths
Pass09:2Ranking	Creates new ranking positions within pattern type lengths
Pass09:3Ranking	Creates relative frequency ranking positions

Appendix 4 - Entity Relationship Diagram

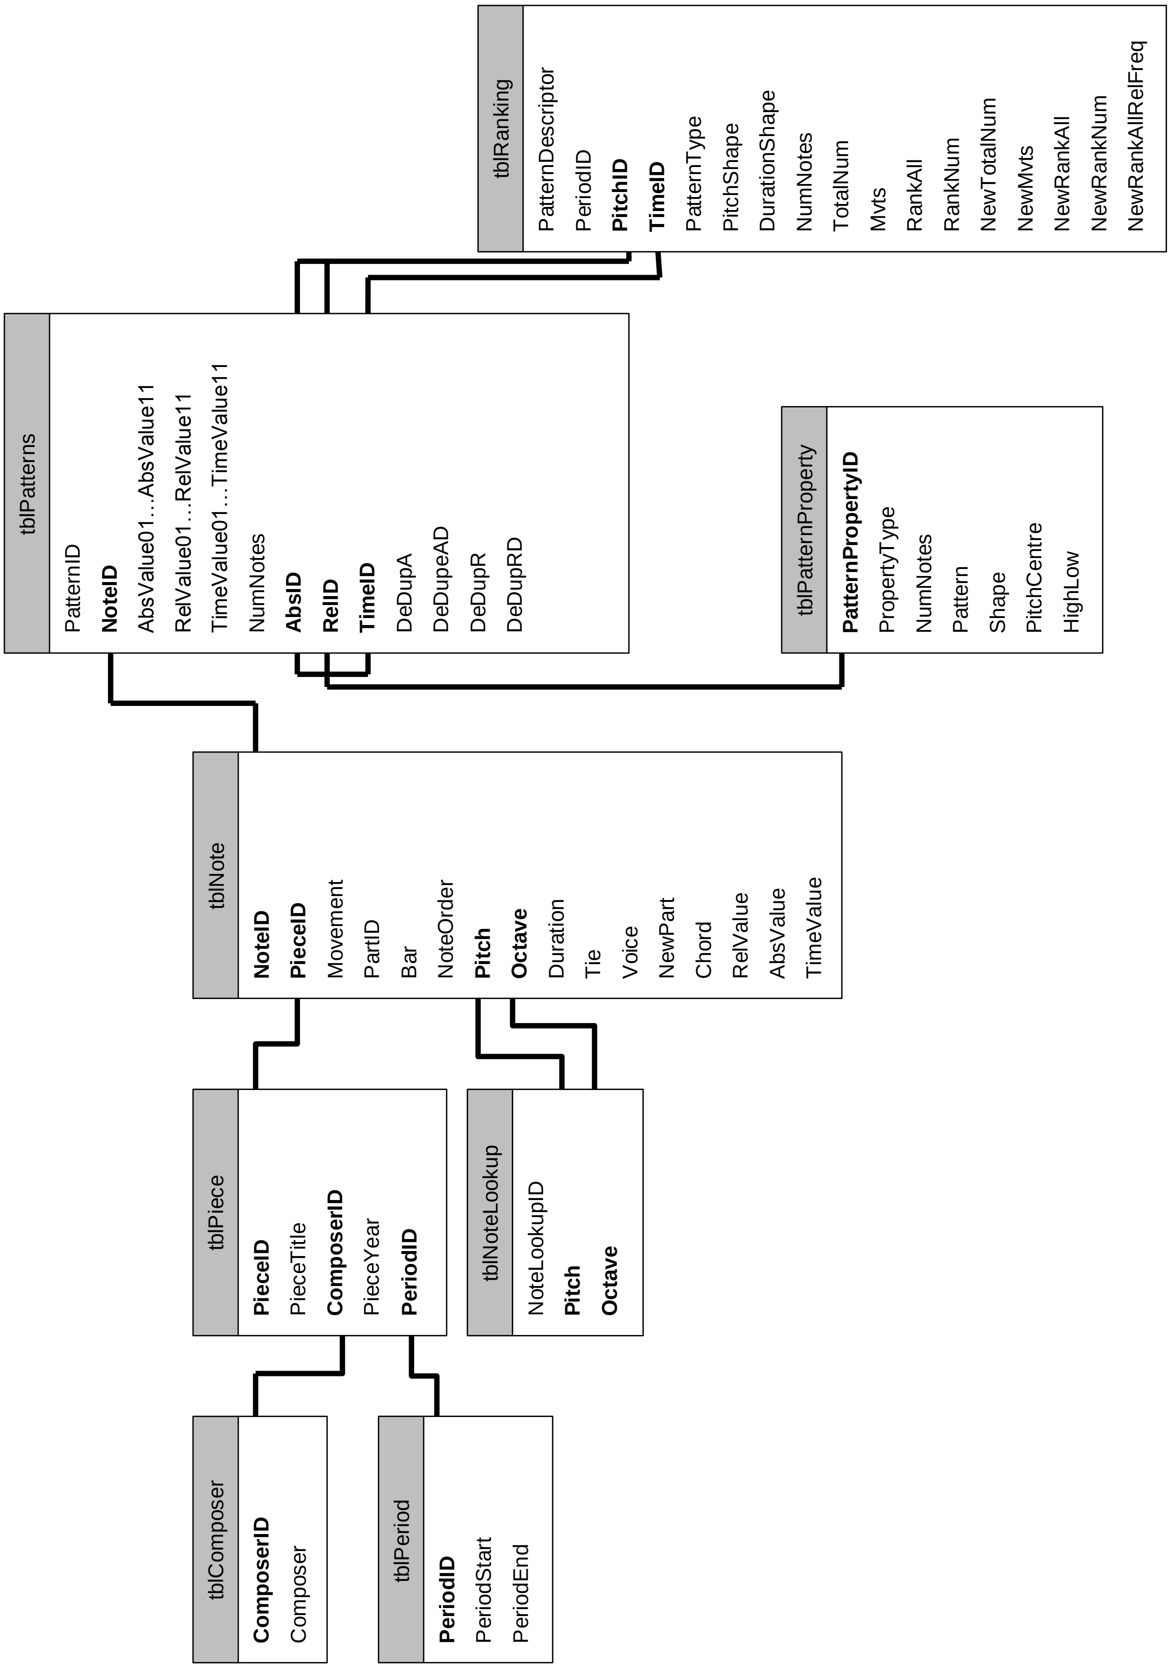

Bibliography

Abercrombie, N., Hill, S. & Turner, B., 2006. Culture. In: N. Abercrombie, S. Hill & B. Turner, eds. *Dictionary of Sociology.* 5th ed. London: Penguin, pp. 92-92.

Adorno, T. W., 1991. *The Culture Industry: Selected Essays on Mass Culture.* London: Routledge.

Agarwal, B. B. & Tayal, S. P., 2009. *Data Mining and Data Warehousing.* New Delhi: University Science Press.

Aunger, R., 2000. Conclusion. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 205-232.

Aunger, R., 2000. Introduction. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 1-24.

Aunger, R., 2002. *The Electric Meme: A New Theory of How We Think.* New York: Free Press.

Baddeley, A., Eysenck, M. W. & Anderson, M. C., 2009. *Memory.* Hove: Psychology Press.

Baddeley, A., Eysenck, M. W. & Anderson, M. C., 2009. *Memory.* Hove: Psychology Press.

Bainbridge, D., 1997. Csound. In: E. Selfridge-Field, ed. *Beyond Midi: The Handbook of Musical Codes.* London: MIT Press, pp. 111-142.

Baragwanath, N., 2011. *The Italian Traditions and Puccini: Compositional Theory and Practice in Nineteenth Century Oper.* Bloomington: Indiana University Press.

Benedict, R., 1934, repr. 1989. *Patterns of Culture.* Boston: Houghton Mifflin.

Best, M. L., 1997. Models for Interacting Populations of Memes: Competition and Niche Behavior. *Journal of Memetics,* 1(2).

Bigand, E., Gérard, Y. & Molin, P., 2009. The Contribution of Local Features to Familiarity Judgments in Music. *Annals of the New York Academy of Sciences,* Volume 1169, pp. 234-244.

Biles, J. A., 2007. Evolutionary Computation for Musical Tasks. In: E. R. Miranda & J. A. Biles, eds. *Evolutionary Computer Music.* London: Springer, pp. 28-51.

Blackmore, S., 1998. Imitation and the Definition of a Meme. *Journal of Memetics,* 2(2).

Blackmore, S., 1999. *The Meme Machine.* New York: OUP.

Blackmore, S., 2000. The Memes' Eve View. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 25-42.

Bloch, M., 2000. A Well-Disposed Social Anthropologist's Problem with Memes. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 189-204.

Bloom, H., 1973. *The Anxiety of Influence: A Theory of Poetry.* New York: OUP.

Blute, M., 2005. Memetics and Evolutionary Social Science. *Journal of Memetics,* 9(1).

Boas, F., 1940. *Race, Language and Culture.* n.p.: s.n.

Borenstein, E., 2004. Survival of the Catchiest: Memes and Postmodern Russia. *The Slavic and East European Journal,* 48(3), pp. 462-483.

Bose, N. K. & Liang, P., 1996. *Neural Network Fundamentals with Graphs, Algorithms, and Applications.* s.l.:McGraw-Hill.

Boyd, R. & Richerson, P. J., 2000. Memes: Universal Acid or a Better Mousetrap?. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 143-162.

Boyle, R. & Coltheart, V., 1996. Effects of Irrelevant Sounds on Phonological Coding in Reading Comprehension and Short-Term Memory. *The Quarterly Journal of Experimental Pshychology,* 49a(2), pp. 398-416.

Brodie, R., 1996. *Virus of the Mind: The New Science of the Meme.* Seattle: Integral Press.

Brower, C., 1993. Memory and the Perception of Rhythm. *Music Theory Spectrum,* 15(1), pp. 19-35.

Brown, A. R., 2004. An Aesthetic Comparison of Rule-Based and Genetic Algorithms for Generating Melodies. *Organised Sound,* 9(2), pp. 191-197.

Burg, J., 2009. *The Science of Digital Media.* London: Pearson Education.

Burkholder, J. P., 1993. Music Theory and Musicology. *The Journal of Musicology,* 11(1), pp. 11-23.

Burkholder, J. P., 1994. The Uses of Existing Music: Musical Borrowing as a Field. *Notes,* 50(3), pp. 851-870.

Butler, M., 2003. Taking it Seriously: Intertextuality and Authenticity in Two Covers by the Pet Shop Boys. *Popular Music,* 22(1), pp. 1-19.

Calvin, W. H., 1996. *How Brains Think: Evolving Intelligence, Then and Now.* New York: Basic Books.

Calvin, W. H., 1997. The Six Essentials? Minimal Requirements for the Darwinian Bootstrapping of Quality. *Journal of Memetics,* 1(1).

Cavalli-Sforza, L. L. & Feldman, M. W., 1981. *Cultural Transmission and Evolution: A Quantitative Approach.* Princeton: Princeton University Press.

Center for Computer Assisted Research in the Humanities, n.d. *Kern Scores.* [Online]
Available at: http://kern.ccarh.org/
[Accessed 2010].

Chen, G., Wu, X. & Zhu, X., 2009. Mining Sequential Patterns across Time Sequences. *New Generation Computing,* Volume 26, pp. 75-96.

Clifford, R., Groult, R., Iliopoulos, C. S. & Byrd, D., 2005. Music Retrieval Algorithms for the Lead Sheet Problem. *Journal of New Music Research,* 34(4), pp. 311-317.

Conklin, D. & Anagnostopoulou, C., 2006. Segmental Pattern Discovery in Music. *INFORMS Journal on Computing,* 16(3), pp. 285-293.

Conklin, D. & Bergeron, M., 2008. Feature Set Patterns in Music. *Computer Music Journal,* 32(1), pp. 60-70.

Conklin, D. & Witten, I. H., 1995. Multiple Viewpoint Systems for Music Prediction. *Journal of New Music Research,* 24(1), pp. 51-73.

Conte, R., 2000. Memes Through (Social) Minds. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 83-120.

Cope, D., 2003. Computer Analysis of Musical Allusions. *Computer Music Journal,* 27(1), pp. 11-28.

Cope, D., 2004. A Musical Learning Algorithm. *Computer Music Journal,* 28(3), pp. 12-27.

Cope, D., 2009. *Hidden Structure: Music Analysis Using Computers.* Middleton: A-R Editions.

Darwin, C., 1859 repr. 1985. *On the Origin of Species by means of Natural Selection.* London: Penguin.

Darwin, C., 1871 repr. 1981. *The Descent of Man and Selection in Relation to Sex.* Princton: Princton University Press.

Darwin, C., 1871, repr. 2009. *The Descent of Man, and Selection in Relation to Sex Volume 1.* Cambridge: Cambridge University Press.

Darwin, C., 1872, repr. 1965. *The Expression of the Emotions in Man and Animals.* Chicago: University of Chicago.

Dawkins, R., 1982. *The Extended Phenotype: The Gene as the Unit of Selection.* Oxford: OUP.

Dawkins, R., 1985. Universal Darwinism. In: D. S. Bendall, ed. *Evolution from Molecules to Men.* Cambridge: Cambridge University Press, pp. 403-428.

Dawkins, R., 1989. *The Selfish Gene.* 2nd ed. Oxford: OUP.

Deitel, P. J. & Deitel, H. M., 2009. *Visual Basic 2008: How to Program.* New Jersey: Prentice Hall.

Dennett, D., 1993. *Consciousness Explained.* London: Penguin.

Dennett, D., 1999. Memes and the Exploitation of Imagination. *The Journal of Aesthetics and Art Criticism,* 48(2), pp. 127-135.

Dennett, D. C., 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life.* London: Penguin Books.

Dirlam, D. K., 2005. Using Memetics to Grow Memetics. *Journal of Memetics,* 9(1).

Dobzhansky, T., 1973. Nothing in Biology Makes Sense Except in the Light of Evolution. *American Biology Teacher,* Volume 35, pp. 125-129.

Dougherty, S., 2001. Culture in the Disk Drive: Computationalism, Memetics, and the Rise of Posthumanism. *Diacritics,* 31(4), pp. 85-102.

Downie, J. S., 2004. A Sample of Music Information Retrieval Approaches. *Journal of the American Society for Information Science and Technology,* 55(12), pp. 1033-1036.

Dunbar, R., 2010. *How Many Friends Does One Person Need?: Dunbar's Number and other Evolutionary Quirks.* London: Faber and Faber.

Durham, W. H., 1991. *Coevolution: Genes, Culture, and Human Diversity.* Stanford: Stanford University Press.

Edmonds, B., 1998. On Modelling in Memetics. *Journal of Memetics,* 2(2).

Edmonds, B., 2002. Three Challenges for the Survival of Memetics. *Journal of Memetics,* 6(2).

Edmonds, B., 2005. The revealed poverty of the gene-meme analogy – why memetics per se has failed to produce substantive results. *Journal of Memetics,* 9(1).

Eliot, T. S., 1949. *Notes Towards the Definition of Culture.* U.S.A.: Harcourt.

Ellen, R., 2002. Dangerous Fictions and Degrees of Plausibility, Evolutionism and Anthropology. *Anthropology Today,* October, 18(5), pp. 3-8.

Elmasri, R. & Navathe, S. B., 2007. *Fundamentals of Database Systems.* 5th ed. Boston: Addison Wesley.

Erikson, R., 1968. Music Analysis and the Computer. *Journal of Music Theory,* 12(2), pp. 240-263.

Forte, A., 1973. The Basic Interval Patterns. *Journal of Music Theory,* 17(2), pp. 234-272.

Forte, A., 1988. New Approaches to the Linear Analysis of Music. *Journal of the American Musicological Society,* 41(2), pp. 315-348.

Gabora, L., 1997. The Origin and Evolution of Culture and Creativity. *Journal of Memetics,* Issue 1.

Gatherer, D., 2005. Finding a Niche for Memetics in the 21st Century. *Journal of Memetics,* 9(1).

Giger, A., 1994. A Bibliography on Musical Borrowing. *Notes,* 50(3), pp. 871-874.

Gimbel, A., 1989. Elgar's Prize Song: Quotation and Allusion in the Second Symphony. *19th-Century Music,* 12(3), pp. 231-240.

Goehr, L., 1994. *The Imaginary Museum of Musical Works: An Essay in the Philosophy of Music.* Oxford: OUP.

Gregg, V. H., 1986. *Introduction to Human Memory.* London: Routledge.

Han, J. & Kamber, M., 2006. *Data Mining: Concepts and Techniques.* 2nd ed. London: Morgan Kaufmann.

Hanninen, D. A., 2001. Orientations, Criteria, Segments: A General Theory of Segmentation for Music Analysis. *Journal of Music Theory,* 45(2), pp. 345-433.

Harris, M., 1997. *Culture, People, Nature: An Introduction to General Anthropology.* 7th ed. New York: Longman.

Haus, G. a. L. M., 2005. A Multi-Layered, Time-Based Music Description Approach Based on XML. *Computer Music Journal,* 29(1), pp. 70-85.

Haus, G., Longari, M. & Pollastri, E., 2004. A Score-Driven Approach to Music Information Retrieval. *Journal of the American Society for Information Science and Technology,* 55(12), pp. 1045-1052.

Haveliwala, T., 2003. Topic-sensitive PageRank: a context-sensitive ranking algorithm for Web search. *Knowledge and Data Engineering, IEEE Transactions on,* 15(4), pp. 784-796.

Haykin, S., 1999. *Neural Networks: A Comprehensive Foundation.* 2nd ed. Upper Saddle River: Prentice Hall.

Hull, D., 2000. Taking Memetics Seriously: Memetics will be what we make it. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 43-68.

Huron, D., 1994. *UNIX Tools for Music Research: The Humdrum Toolkit; Reference manual.* s.l.:Center for Computer Assisted Research in the Humanities.

Huron, D., 1997. Humdrum and Kern: Selective Feature Encoding. In: E. Selfridge-Field, ed. *Beyond Midi: The Handbook of Musical Codes.* London: MIT Press, pp. 375-401.

Huron, D., 1998. *Music Research using Humdrum: A User's Guide.* [Online]
Available at: http://www.musiccog.ohio-state.edu/Humdrum/guide.toc.html
[Accessed 22 February 2010].

Huron, D., 2002. Music Information Processing using the Humdrum Toolkit: Concepts, Examples, and Lessons. *Computer Music Journal,* 26(2), pp. 11-26.

Huron, D., 2006. *Sweet Anticipation: Music and the Psychology of Expectation.* Massachusetts: MIT Press.

Husbands, P., Copley, P., Eldridge, A. & Mandelis, J., 2007. An Introduction to Evolutionary Computing for Musicians. In: E. R. Miranda & J. A. Biles, eds. *Evolutionary Computer Music.* London: Springer, pp. 1-27.

International Music Score Library Project, n.d. *Petrucci Music Library.* [Online]
Available at: http://imslp.org/wiki/Main_Page
[Accessed 2010].

Jackendoff, R. & Lerdahl, F., 2006. The Capacity for Music: What is it, and What's Special About It?. *Cognition,* 100(1), pp. 33-72.

Jan, S., 2000. Replicating Sonorities: Towards a Memetics of Music. *Journal of Memetics,* 4(1).

Jan, S., 2003. The Evolution of a 'Memeplex' in Late Mozart: Replicated Structures in Pamina's 'Ach ich fühl's. *Journal of the Royal Musical Association,* Volume 128, pp. 330-370.

Jan, S., 2007. *The Memetics of Music: A Neo-Darwinian View of Musical Structure and Culture.* Aldershot: Ashgate Publishing Limited.

Jan, S., 2010. Memesatz contra Ursatz: Memetic perspectives on the aetiology and evolution of musical structure. *Musicae Scientiae,* 14(1), pp. 3-50.

Jan, S., 2011. Music, Memory, and Memes in the Light of Calvinian Neuroscience. *Music Theory Online,* 17(2).

Jan, S., 2012. 'The Heavens are Telling': Memetic- Calvinian Readings of a Haydn Chord Progression. *Interdisciplinary Science Review,* 37(2), pp. 113-130.

Jan, S., 2013. Similarity Continua and Criteria in Memetic Theory and Analysis. *Under Review.*

Kaplan, D., 1960. The Law of Cultural Dominance. In: M. D. Sahlins & E. R. Service, eds. *Evolution and Culture.* Michigan: University of Michigan Press, pp. 69-92.

Klein, M. L., 2005. *Intertextuality in Western Art Music.* Bloomington: Indiana University Press.

Knapp, R., 1998. Brahms and the Anxiety of Allusion. *Journal of Musicological Research,* 18(1), pp. 1-30.

Kolb, B. & Whisha, I. Q., 1998. Brain Plasticity and Behavior. *Annual Review of Psychology,* Volume 49, pp. 43-64.

Korsyn, K., 1991. Towards a New Poetics of Musical Influence. *Music Analysis,* 10(1/2), pp. 3-72.

Kostek, B., 2004. Application of Soft Computing to Automatic Music Information Retrieval. *Journal of the American Society for Information Science and Technology,* 55(12), pp. 1108-1116.

Koutsobina, V., 2008. Readings of Poetry – Readings of Music: Intertextuality in Josquin's Je me Complains de Mon Amy. *Early Music,* 36(1), pp. 67-78.

Kroeber, A., 1952. *The Nature of Culture.* Chicago: Chicago University Press.

Krumhansl, C. L., 1995. Music Psychology and Music Theory: Problems and Prospects. *Music Theory Spectrum,* 17(1), pp. 53-80.

Krumhansl, C. L., 2004. The Cognition of Tonality – As we Know it Today. *Journal of New Music Research,* 33(3).

Kuper, A., 2000. If Memes are the Answer, then what is the Question?. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 175-188.

Laland, K. & Odling-Smee, J., 2000. The Evolution of the Meme. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 121-142.

Langrish, J. Z., 1999. Different Types of Memes: Recipemes, Selectemes and Explanemes. *Journal of Memetics,* 3(2).

Lartillot, O., 2004. A Musical Pattern Discovery System founded on a Modeling of Listening Strategies. *Computer Music Journal,* 28(3), pp. 53-67.

Lartillot, O., 2005. Multi-Dimensional Motivic Pattern Extraction founded on Adaptive Redundancy Filtering. *Journal of New Music Research,* 34(4), pp. 375-393.

Lefkowitz, D. S. & Taavola, K., 2000. Segmentation in Music: Generalizing a Piece-Sensitive Approach. *Journal of Music Theory,* 44(1), pp. 171-229.

Levi-Strauss, C., 1962. *La Pensée Sauvage.* Paris: Plon.

Linderholm, C. E., 1972. *Mathematics Made Difficult.* New York: World Publishing.

Lynch, A., 1998. *Thought Contagion: How Belief Spreads Through Society.* New York: Basic Books.

Maimon, O. Z. & Rokach, L., 2005. *Data Mining and Knowledge Discovery Handbook.* USA: Springer.

Mason, S. & Saffle, M., 1994. L-Systems, Melodies and Musical Structure. *Leonardo Music Journal,* Volume 4, pp. 31-38.

Melucci, M. & Orio, N., 2004. Combining Melody Processing and Information Retrieval Techniques: Methodology, Evaluation, and System Implementation. *Journal of the American Society for Information Science and Technology,* 55(12), pp. 1058-1066.

Meredith, D., Lemström, K. & Wiggins, G. A., 2002. Algorithms for Discovering Repeated Patterns in Multidimensional Representations of Polyphonic Music. *Journal of New Music Research,* 31(4), pp. 321-345.

Metzer, D., 2003. *Quotation and Cultural Meaning in Twentieth-Century Music.* Cambridge: Cambridge University Press.

Micznik, V., 2000. Of Ways of Telling, Intertextuality, and Historical Evidence in Belioz's Roméo et Juliette. *19th-Century Music,* 24(1), pp. 21-61.

Milicevic, M., 1998. Deconstructing Musical Structure. *Organised Sound,* 3(1), pp. 27-34.

Milicevic, M., 2001. Cyberspace Memes. *Organised Sound,* 6(2), pp. 117-120.

Miller, G. A., 1956. The Magical Number Seven, Plus or Minus Two: Some Limits on our Capacity for Processing Information. *The Psychological Review,* Volume 63, pp. 81-97.

Miranda, E. R., Kirby, S. & Todd, P., 2003. On Computational Models of the Evolution of Music: From the Origins of Musical Taste to the Emergence of Grammars. *Contemporary Music Review,* 22(2), pp. 91-110.

Miranda, E. R. & Todd, P. M., 2007. Computational Evolutionary Musicology. In: E. R. Miranda & J. A. Biles, eds. *Evolutionary Computer Music.* London: Springer, pp. 28-51.

Morris, R. D., 1993. New Directions in the Theory and Analysis of Music Contour. *Music Theory Spectrum,* 15(2), pp. 205-228.

Narmour, E., 1990. *The Analysis and Cognition of Basic Melodic Structures: The Implication-Realization Model.* Chigaco: University of Chigaco Press.

Parmer, D., 1995. Brahms, Song Quotation, and Secret Programs. *19th Century Music,* 19(2), pp. 161-190.

Patterson, C., 1987. Evolution: Neo-Darwinism. In: R. L. Gregory, ed. *The Oxford Companion to the Mind.* Oxford: OUP, pp. 234-244.

Pearsall, E., 1999. Mind and Music: On Intentionality, Music Theory, and Analysis. *Journal of Music Theory,* 43(2), pp. 231-255.

Peraino, J. A., 2001. Re-Placing Medieval Music. *Journal of the American Musicological Society,* 54(2), pp. 209-264.

Plasketes, G., 2005. Re-flections on the Cover Age: A Collage of Continuous Coverage in Popular Music. *Popular Music and Society,* 28(2), pp. 137-161.

Plotkin, H., 2000. Cultural and Psychological Mechanisms. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 69-82.

Plotkin, H. C., 1994. *Darwin Machines and the Nature of Knowledge.* Cambridge: Harvard University Press.

Polansky, L., 1996. Morphological Metrics. *Journal for New Music Research,* 25(4), pp. 289-368.

Pollack, R. D., 2008. Data Mining: Common Definitions, Applications, and Misunderstandings. In: K. D. Lawrence, S. Kudyba & R. K. Klimberg, eds. *Data Mining Methods and Applications.* London: Auerbach, pp. 229-238.

Pople, A., 2002. *Getting Started with the Tonalities Music Analysis Software.* [Online]
Available at:
http://www.nottingham.ac.uk/shared/shared_music/documents/Getting_Started_with_Tonalities.pdf
[Accessed 10 December 2012].

Prigmore, M., 2008. *An Introduction to Databases with Web Applications.* Harlow: Pearson Education.

Project Gutenberg, n.d. *Project Gutenberg: The Sheet Music Project.* [Online]
Available at: http://www.gutenberg.org/wiki/Gutenberg:The_Sheet_Music_Project
[Accessed November 2010].

Rabuñal, J. R. & Dorado, J., 2008. *Artificial Neural Networks in Real-Life Applications.* London: Idea Group Publishing.

Radvansky, G., 2006. *Human Memory.* Boston: Pearson.

Recanati, F., 2001. Open Quotation. *Mind,* 110(439), pp. 637-687.

Recordare, 2008. [Online]
Available at: http://www.recordare.com/xml/musicxml-tutorial.pdf
[Accessed 22 February 2010].

Rizo, D., Iñesta, J. M. & Lemström, K., 2011. Polyphonic Music Retrieval with Classifier Ensembles. *Journal of New Music Research,* 40(4), pp. 313-324.

Rolland, P.-Y., 1999. Discovering Patterns in Musical Sequences. *Journal of New Music Research,* 28(4), pp. 334-350.

Rose, N., 1998. Controversies in Meme Theory. *Journal of Memetics,* 2(1).

Sahlins, M. D., 1960. Evolution: Specific and General. In: M. D. Sahlins & E. R. Service, eds. *Evolution and Culture.* Michigan: University of Michigan Press, pp. 2-44.

Sanguinetti, G., 2007. The Realization of Partimenti: An Introduction. *Journal of Music Theory,* 51(1), pp. 51-83.

Santana, H., Dahia, M., Lima, E. & Ramalho, G., 1998. *VExPat: An Analysis Tool for the Discovery of Musical Patterns.* [Online]
Available at: http://en.scientificcommons.org/43555950
[Accessed 23 02 2010].

Selfridge-Field, E., 1997. Introduction: Describing Musical Information. In: E. Selfridge-Field, ed. *Beyond Midi: The Handbook of Musical Codes.* London: MIT Press, pp. 2-40.

Sloan, D., 1997. HyTime and Standard Music Description Language: A Document-Description Approach. In: E. Selfridge-Field, ed. *Beyond Midi: The Handbook of Musical Codes.* London: MIT Press, pp. 469-590.

Sloboda, J., 2005. *Exploring the Musical Mind: Cognition, Emotion, Ability, Function.* Oxford: OUP.

Sloboda, J. A., 1985. *The Musical Mind: The Cognitive Psychology of Music.* Oxford: Clarendon Press.

Snyder, B., 2000. *Music and Memory: An Introduction.* Cambridge: MIT Press.

Sotiropoulos, D. N., Lampropoulos, A. S. & Tsihrintzis, G. A., 2008. MUSIPER: A System for Modeling Music Similarity Perception based on Objective Feature Subset Selection. *User Modeling and User-Adapted Interaction,* Volume 18, pp. 315-348.

Sperber, D., 2000. An Objection to the Memetic Approach to Culture. In: R. Aunger, ed. *Darwinizing Culture: The Status of Memetics as a Science.* Oxford: OUP, pp. 163-174.

Sterelny, K., 2006. Memes Revisited. *British Society for the Philosophy of Science,* 57(1), pp. 145-165.

Steward, J., 1955. *Theory of Culture Change.* Urbana: Illinois University Press.

Stewart, D., 2003. XML for Music: A Markup Language that Breaks Down Musical Barriers. *Electronic Musician*, December, pp. 58-64.

Taniar, D., 2008. *Data Mining and Knowledge Discovery Technologies.* London: IGI Publishing.

Tarnawska-Kaczorowska, K., 1998. Musical Quotation an Outline of the Problem. *Contemporary Music Review,* 17(3), pp. 69-90.

Temperley, D., 2001. *The Cognition of Basic Musical Structures.* Cambridge, Mass.: MIT Press.

Temperley, D., 2007. *Music and Probability.* Cambridge, Mass.: MIT Press.

Uitdenbogerd, A. L. & Zobel, J., 2004. An Architecture for Effective Music Information Retrieval. *Journal of the American Society for Information Science and Technology,* 55(12), pp. 1053-1057.

Vaneechoutte, M. & Skoyles, J. R., 1998. The Memetic Origin of Language: Modern Humans as Musical Primates. *Journal of Memetics,* 2(2).

Various, 2011. *Grove Music Online.* [Online]
Available at: http://www.oxfordmusiconline.com
[Accessed 2011].

Verberk, W. C. E. P., 2012. Explaining General Patterns in Species Abundance and Distributions. *Nature Education Knowledge,* 3(10), p. 38.

Volkov, I., Banavar, J. R., Hubbell, S. P. & Maritan, A., 2003. Neutral Theory and Relative Species Abundance in Ecology. *Nature,* 424(6952), p. 19.

Walker, R., 2004. Cultural Memes, Innate Proclivities and Musical Behaviour: A Case Study of the Western Traditions. *Psychology of Music,* April, 21(2), pp. 153-190.

Wheatland, P., n.d. *Melodic Match: Music Analysis Software.* [Online]
Available at: http://www.melodicmatch.com/
[Accessed 22 02 2010].

White, B. W., 1960. Recognition of Distorted Melodies. *American Journal of Psychology,* 73(1), pp. 100-107.

White, L., 1959. *The Evolution of Culture: The Development of Civilization to the Fall of Rome.* New York: McGraw-Hill.

Wilkins, J. S., 1998. What's in a Meme? Reflections from the Perspective of the History and Philosophy of Evolutionary Biology. *Journal of Memetics,* 2(1).

Williams, A., 2006. Review: Michael L. Klien: Intertextuality in Western Art Music. *Music Theory Spectrum,* 28(1), pp. 316-320.

Wilson, C., 2009. *Charles Darwin Tagged You in a Note on Facebook: The evolutionary roots of Facebook's "25 Things" craze.* [Online]
Available at: http://www.slate.com/id/2211068/
[Accessed 14 February 2009].

Witten, I. H. & Frank, E., 2000. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations.* San Francisco: Morgan Kaufmann.

Wong, Y.-L., 2009. *Digital Media Primer: Digital Audio, Video, Imaging and Multimedia Programming.* London: Prentice Hall.

Xiao, Q., Tsuge, S. & Kita, K., 2011. *Music Retrieval Method Based on Filter-bank Feature and Earth Mover's Distance.* Shanghai, IEEE, pp. 1845 - 1849.

Zack, E. S., 2005. Theorising Music History: Intertextuality as a Tool for Reclaiming the Past in Fanny Mendelssohn's Canata Lobgesang. *Women's Writing,* 12(1), pp. 59-72.