How can music visualisation techniques reveal different perspectives on musical structure?

Samuel J. Hunt
Department of Computer Science & Creative Technologies.
University of The West of England
`Samuel.hunt@uwe.ac.uk`

Tom Mitchell
Department of Computer Science & Creative Technologies.
University of The West of England
`Tom.mitchell@uwe.ac.uk`

Chris Nash
Department of Computer Science & Creative Technologies.
University of The West of England
`Chris.nash@uwe.ac.uk`

ABSTRACT

Standard western notation supports the understanding and performance of music, but has limited provisions for revealing overall musical characteristics and structure. This paper presents several visualisers for highlighting and providing insights into musical structures, including rhythm, pitch, and interval transitions, also noting how these elements modulate over time. The visualisations are presented in the context of Shneiderman's Visual Information-Seeking Mantra, and terminology from the Cognitive Dimensions of Music Notations usability framework. Such techniques are designed to make understanding musical structure quicker, easier, less error prone, and take better advantage of the intrinsic pattern recognition abilities of humans.

1. INTRODUCTION

Standard western notation serves as a strict, formal set of instructions for the performance of composed music. However, it omits explicit representation of a rich amount of hidden data that exists between individual notes, and the location of the notes within an overarching musical structure. One way to understand this structure is to analyse the music: either manually, requiring an experienced musicologist; or via computer, resulting in several multi-dimensional data fields, which may be difficult to represent and comprehend. Representing this data visually utilises the brain's pattern detection abilities, supporting easier and faster comprehension of material to enable insight and speculation that can inform further formal analysis.

Visualisation presents non-visual data in a visual format, usually as 2D/3D images or video. Shneiderman [1] introduces a framework for guiding the design of information visualisation systems, known as the Visual Information-Seeking Mantra (VISM). The framework consists of seven tasks for presenting information in a visual form to a user (Table 1). Craft and Cairns [2] elaborate on this by stating the VISM serves as inspiration and guidelines for practitioners designing visual information systems.

Task	Description
Overview	Gain an overview of the data.
Zoom	Zoom in on items of interest.
Filter	Filter out uninteresting items.
Details-on-Demands	Selected an item or group and get details when needed.
Relate	View relationships between items.
History	Keep a history of actions to support undo, replay, and progressive refinement.
Extract	Allow extraction of sub-collections and of the query parameters.

Table 1. The 7 tasks of the VISM.

Shneiderman emphasises that humans have remarkable perceptual abilities, allowing them to easily detect changes of and patterns in size, colour, shape, movement or texture in visual media. Such advanced and robust feature extraction capabilities are considerably more difficult to encode as automated analysis using computer systems.

In a musical context, visualisers also enable rapid, automated methods for visualising not only a single piece of music, but an entire corpus - allowing understanding and comparisons of musical material at a higher and more generalised level to that of manual score analysis.

The level meter which features in the majority of consumer audio products, represents a ubiquitous visualisation method, whereby the current sound level is visualised using vertical bars, and for the majority of situations a more useful presentation than a display of audio sample values (amplitudes). Digital audio is stored as a series of numbers, a sequence of amplitude measurements with respect to time. Sonograms convert this information to visualise the distribution of frequency content. An example of this is illustrated in Figure 1, whereby the musical score has been synthesized using piano samples on a computer and analysed with a sonogram.

This paper focuses on visualizing scores at the note-level (e.g. MIDI), avoiding the many difficulties of audio feature extraction. Sequenced music, encoded as MIDI, by contrast allows for rapid and reproducible analysis [3]. The aim of the paper is to present novel techniques that support the analysis of music.
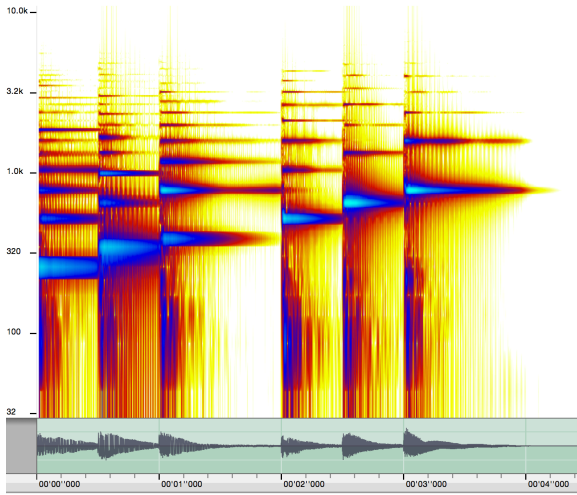
Figure 1. Sonogram plot of the score.

Figure 2. Self-Similarity visualisation of the score

The remainder of the paper is broken down as follows. Section 2 presents relevant prior work and theory, followed by a brief discussion in Section 3 of the software system developed to support this research. Section 4 reviews visualisation techniques for pitch, contours, intervals and key, followed by Section 5 looking at rhythmic elements. Section 6 discusses visualisation techniques that integrate both pitch and rhythmic elements. The final section considers future work for the area and proposes evaluation techniques.

2. RELATED WORK

Prior work in music visualisation can be broadly categorised into two groups: those exploring sampled audio data and those exploring sequenced music data (scores and MIDI). Soriano et al [4] present methods for browsing an audio-based music collection, using graphical metaphors designed to convey the underlying song structure. This analysis is performed via feature extraction from MIDI files, enabling easy identification of simple and meaningful musical structure, such as pitch and rhythm.

Foote [5] and Wolkowicz & Brooks [6] both used self-similarity matrix visualisations to reveal similarity in music. This visualisation approach relies on the measurement of pitch content at quantised time intervals, and plotting this against all other intervals. Figure 2 shows a self-similarity matrix visualisation, whereby the music proceeds through time from the bottom left to the top right, with regions of similar patterns appearing as clusters of squares. Both axes represent the same input vector. The music example uses a repeating motif of one bar, with a modulation at bars 2 and 4.

Bergstrom [7] presents several visualisers that convey information about interval quality, chord quality, and the chord progressions in a piece of music, helping users to comprehend the underlying structure of music.
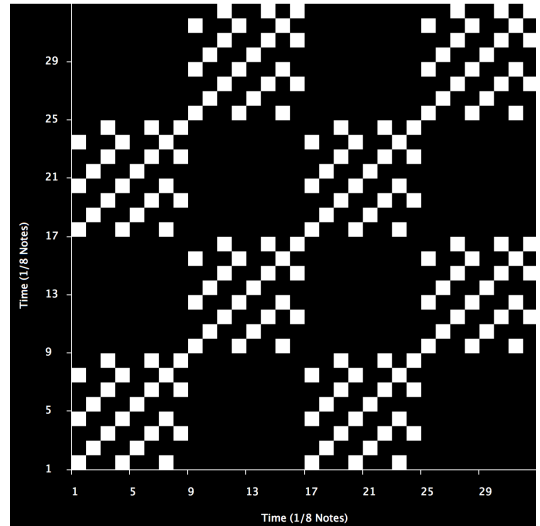
Feedback from engagement with the system revealed users who having quickly understood the basics, wanted to compare music from multiple genres and composers. Holland [8] presents a similar system (Harmony Space) to allow beginners to interact with harmony using a visual grid.

Jeong and Nam [9] discuss a system that visualises audio streams, to show audio features such as, volume, onset density, and dissonance. The authors also state that as music is an auditory art, visual representations can contain information that cannot be transferred or perceived accurately with sound. Herremans and Chew [10] use visualisation to highlight tonal tension in music, creating an explicit representation of something that is not easily quantifiable, presenting graphics alongside the scored elements.

Established analytical frameworks for music, such as the Generative Theory of Tonal Music (GTTM) [11] and Schenkerian analysis [12], also present ways to annotate music and reveal structure. The GTTM proposes a series of preference rules for determining the different musical structures that underlie the perception of western music. Schenkerian analysis is an established musical analysis technique that aims to explicitly reveal hidden dependencies and structures implicit in the music. This analysis primarily aids score reading by marking it with elements of musical structure. Both of these theories have been mechanised in software [13][14].

Nash [4] presents research that adapts the Cognitive Dimensions of Music Notations framework (CDMN) [15], for use in designing and analyzing music notations and user interfaces for digital and traditional music practice and study. This paper utilises the framework as a vocabulary for comparing visualised music content and metadata against western notation and other forms of visualisation. However not all of the 16 core dimensions' originally specified are of relevance here. A list of the terms and their definitions relevant in this research are listed in Table 2.

Dimension	Description
Visibility	How easy is it to view and find elements of the music during editing?
Juxtaposabillity	How easy is it to compare elements within the music?
Hidden Dependencies	How explicit are the relationships between related elements in the notation?
Hard Mental Operations	How difficult is the task to work out in your head?
Conciseness	How concise is the notation?
Provisionality	How easy is it to experiment with ideas?
Consistency	Where aspects of the notation mean similar things, is the similarity clear in the way they appear?
Viscosity	Is it easy to go back and make changes?
Role Expressiveness	Is it easy to see what each part of the notation means?
Error Proneness	How easy is it to make annoying mistakes?
Closeness of mapping.	Does the notation match how you describe the music yourself?

Table 2. Terms of the Cognitive Dimensions of Music Notations framework used in this paper [4].

Using the self-similarity visualisation in Figure 2 as an example of the terms used in the table, the *visibility* of the figure is good, showing a clear overview of the entire piece, likewise the *juaxtaposability* scores highly as the patterns can be compared much more easily than sequences in the score. There are high *hidden dependences* as the original information has been transformed, with each square representing a smaller amount of information. The simplistic nature of the visualiser scores high on *provisionality, consistency and conciseness*. The visualisation does not have any meaning unless related to the score, with the underlying notated elements looked up, so has a poor *closeness of mapping*. Comparing sequences using just the notation would require both hard mental operations, and would be prone to error (*error proness*), whereas the automated analysis used to build the self-similarity visualisation is easily reproducible and more accurate.

A core concept of visualisation for notated music is its ability to reduce the *hard-mental operations* arising from manual score analysis [4]. Computer aided analysis also reduces the *error proneness* of operations. Visualisation can remove un-needed details (*filtering* [1]) from the score, for example performance markings, therefore improving the *conciseness* of the results.

Temperley [16, 17] uses visualisation to inform, explain, and evaluate formal analysis by computer. Often using these techniques when analysing a large corpus of music, to immediately show data that would otherwise be difficult to extract from looking directly at the score, or in fact thousands of individual pieces. Temperley also uses these as a way of comparing and refining models for music analysis.

3. INTERACTIVE VISUALISATION

This paper discusses visualisers developed for an original software package (Figure 3), the design of which has been influenced by the seven principles of the VISM (listed in Table 1). In general, it allows different pieces of music in MIDI format, to be opened and visualised quickly, in order to support *high provisionality* and enable rapid experimentation with analysis techniques. The software can analyse and compare entire corpora or individual pieces, as well as sub-sections or voices (tracks). A historic list of analyses is kept so these can be recalled and modified, retaining low levels of *viscosity* and commitment, therefore further facilitating experimentation and evaluation (*provisionality)*.

Software and automated analysis has the advantage of processing large amounts of data quickly (compared to manual techniques), but takes considerable amounts of time and care to design and implement. Visualisation tools, such as that described, allow a user to speculatively interrogate data, before committing to more detailed and formal music analysis methods, be they traditional (e.g. Schenkerian) or computer-based (e.g. machine learning see Section 4.4).

4. VISUALISING PITCH

The set of visualisers presented in this section focus on elements of pitch, contour, and melodic interval. Some techniques present the material as overviews of the piece as a whole, others present excerpts in time. For the purposes of discussion and comparison, the majority of visualisations present Bach's *Two-part Invention No. 1 (BWV 772)* [18], but can be applied to many other examples and genres of music, including non-Western.

4.1 Melodic Contours

A contour representation of music can simply be defined as information about the up and down pattern of pitch changes, regardless of their exact size [19]. Melodic contours are also a key psychological part of music, one that aids the recollection of musical themes [20].
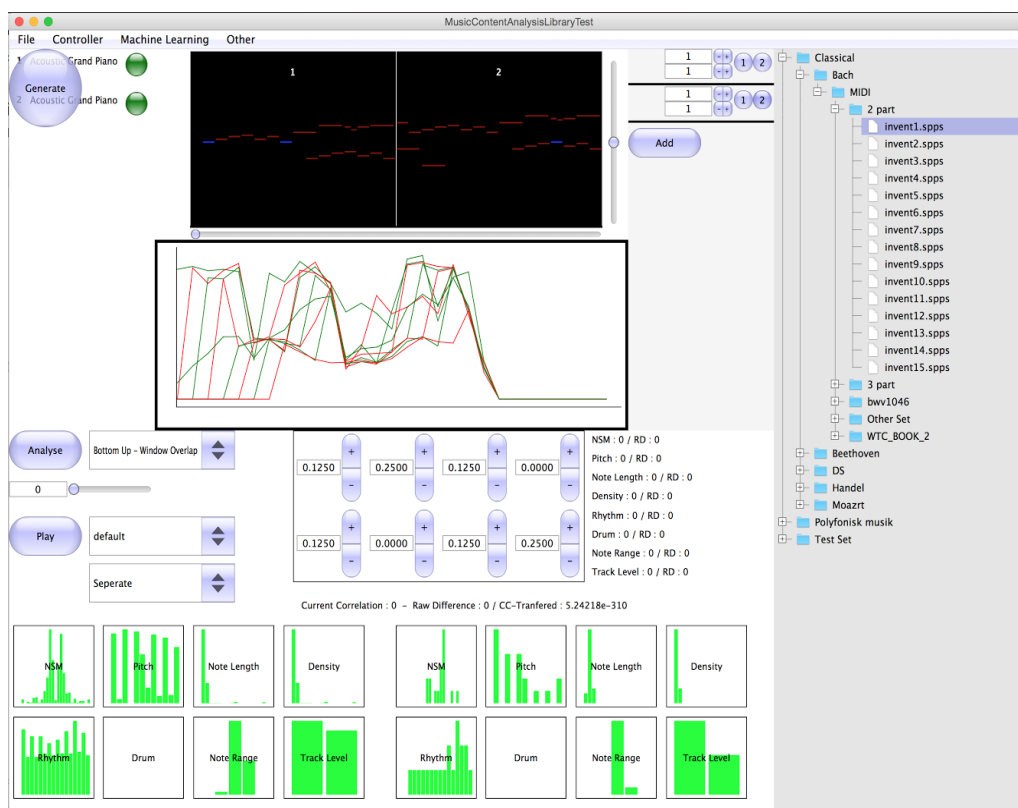
Figure 3. Software created to support visualisation tasks.

Melodic contours themselves can be illustrated using a score, where it is usually clear in which direction the pitch is going (Figure 4, top). However, once accidentals are introduced (Figure 4, bottom), it becomes less visually distinct. A piano roll (Figure 5) provides a clearer representation of melodic contour. This provides improved *closeness of mapping* [2], and increases the ease with which sequences can be compared (improving *juxtaposability*). Piano rolls provide ways for shapes, patterns and contours to be identified. Wood [21] presents related research in which the standard note head is visually modified to show the pitch degree in a more role expressive way, and reports improved speed for sight-reading when compared with standard note heads.

Figure 4. Score with clear melodic contour (top) and obfuscated melodic contour (bottom).

This type of visualisation can also be used to reduce a search space, allowing sequences represented as contours to be visually clustered. The items in Figure 6 show a series of monophonic melodies extracted from Bach's BWV 772. Visually, we can see that the first two patterns

are similar, and that pattern 14 is the same pattern inverted. This kind of visualisation allows the viewer to employ the gestalt principles of visual perception, in this case similarity, to group together similar shapes [22]. In this situation the data has *filtered* out everything but the contour, giving a better *overview* of the types of contours, which can then be easily *related* against one another.
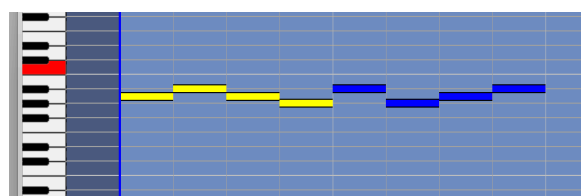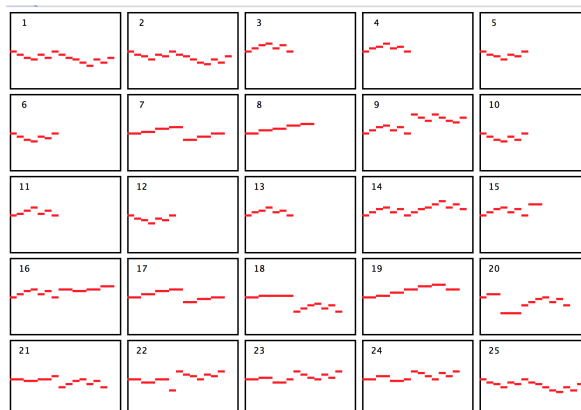

Figure 5. Piano roll representation of Figure 4.

Figure 6. Selection of melodic contours from Bach's BMW 772.

4.2 Intervals

The contour plots provide an overview of the melodic patterns present in the music, but reduce the *visibility* and *role expressivity* of the intervals. Temperley [16, 17] uses a histogram of melodic intervals to show the distribution of interval leaps between melodic note sequences within an entire corpus of music material, revealing wider patterns and trends in music. In-so-doing, this *hides dependencies* in the music, such as the local context and note-to-note relationships (i.e. certain pitches are more unlikely to transition to those depicted in the figure because of their relation to the home key and sensitivity to tonal context). The diagram in Figure 7 shows the interval profile for Bach's BWV 772.
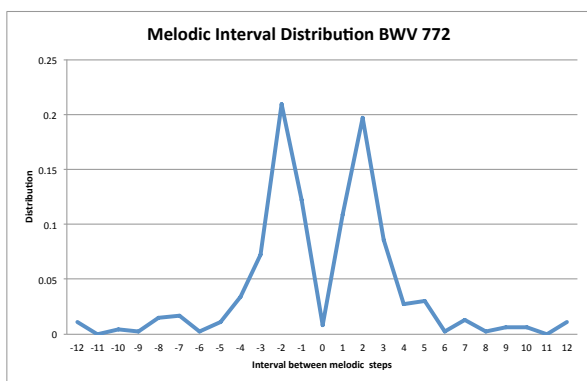

Figure 7. Interval distribution over two octaves in Bach's BWV 772.

A different way to analyse this data, in a way that allows interpretation of pitch, intervals and range, is to use a Markov type model, defining the transition probability between any given notes, in a numeric table format. This, however produces a data table of size 127x127 elements, which is difficult to comprehend in a numeric format, but easily visualised to reveal musical trends and characteristics – as illustrated in Figure 8. The design of this once again takes an *overview* of the data, *filtering* out the timed elements of the music, to give a detailed overview of the pitch and interval elements. Parts of the plot can be further inspected to reveal exact transition probabilities (*details on demand*).

From the plot, it can be noted that the intervals in the upper ranges are more likely to jump down in interval, while the opposite effect can be observed in the lower range. Towards the middle the width of the melodic jumps are slightly larger. The yellow line along the leading diagonal represents the unison interval (repeated notes), the horizontal deviation from which reflects transitions to subsequent notes. The brighter the marker, the more likely the transition. The diagram can also be thought of as a layered series of melodic interval distributions (as in Figure 7), given different starting notes (y-axis).
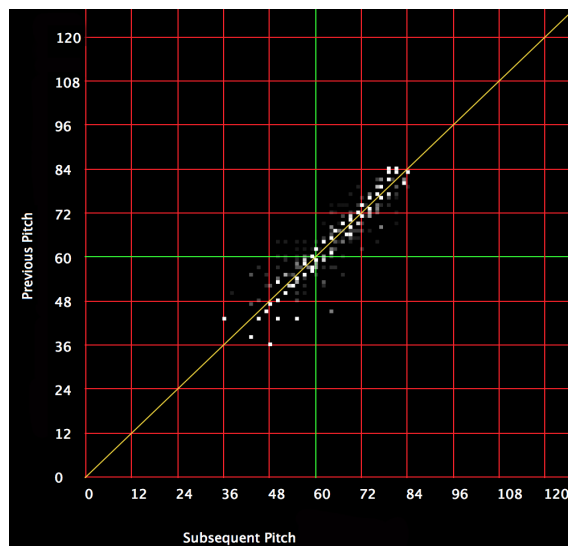

Figure 8. 2D Markov plot of Bach's BWV 772.

Two more plots are shown in this style, but illustrating trends in, and differences between, larger corpora of music: respectively, a collection of 100 pieces of baroque music (Figure 9) and jazz music (Figure 10), selected randomly from a larger corpus. The visualisation process helps to reveal differences between the corpora that would otherwise be harder to discover or articulate. For example, the range of intervals in the jazz corpus is far wider, whereas the baroque is limited to mostly to an octave, and multiples thereof and appears more uniform throughout the range.
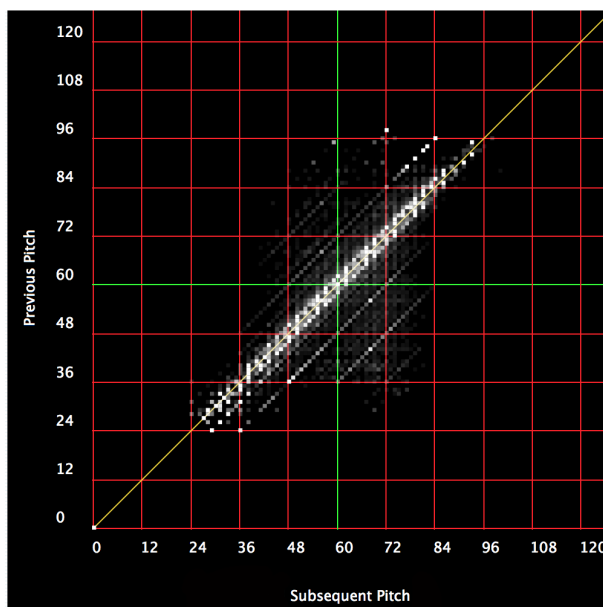

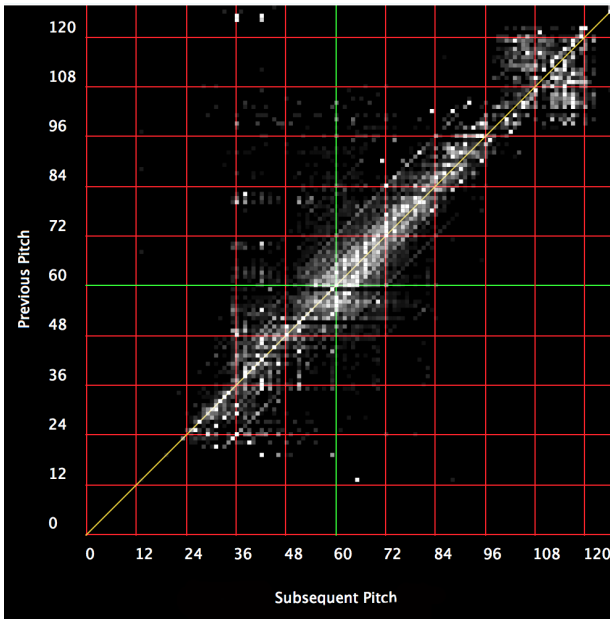**Figure 9.** 100 randomly-selected common repertoire Baroque pieces.

Figure 10. 100 randomly-selected common repertoire Jazz pieces.

4.3 Pitch Distribution

It is instructive to consider pitch usage in general terms. Temperley [16,17] considers the distribution of pitches within a piece to be an intrinsic element that grounds the overall tonality and key in western music. Key is something that musicians are trained to detect [16], but for which Temperley has developed automated methods. To illustrate, Figure 11 shows an ideal key profile describing the average distribution of pitches within a piece in C major, which can also be considered a coarse measure of pitch-class appropriateness in relation to key. For comparison Bach's BWV 772 (Figure 12) is also visualised. It is easy to visually infer the similarity of the distribution within the piece (known to be in the key of C) and the generalised representation (Figure 11). Smaller more nuanced details are also visible, such as the fact that the piece, although in C major, has more instances of D than the tonic C. Such details can be enough to fool automated analysis, as detailed in the next section, but things are clearer to the eye.
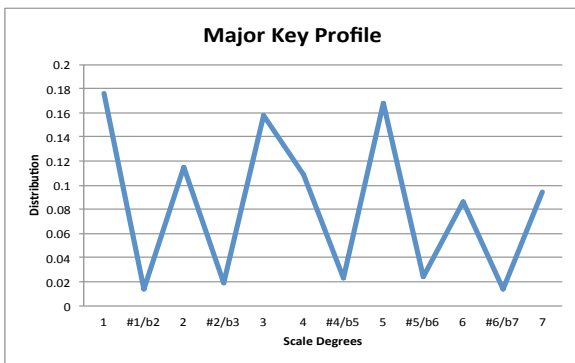

Figure 11. Major Key Profile.

Other metadata can also loosely be inferred. A less pronounced distribution may indicate a piece that uses several different keys or tonalities beyond the diatonic. Atonal music, such as serialism, may confound such analysis and appear entirely different when visualised, such as Schoenberg *Op.11-1* (Figure 13).
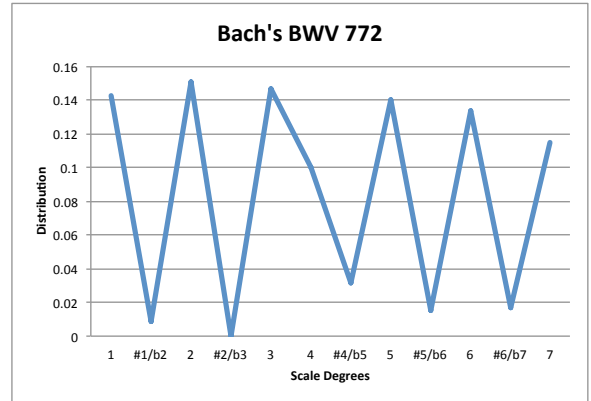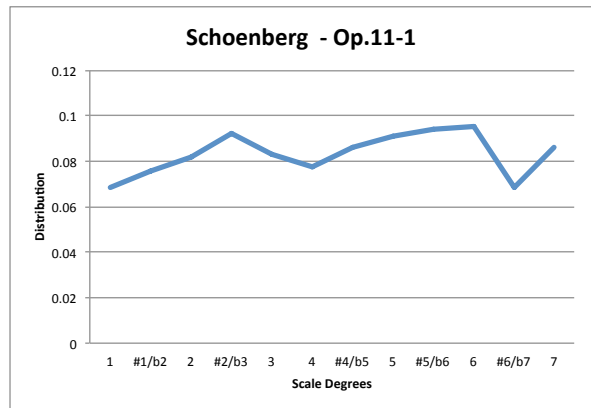

Figure 12. Pitch distribution in Bach's BWV 772

Figure 13. Pitch Distribution in Schoenberg Op.11-1 [24].

4.4 Key

Visualisation can help guide and test formal analysis. For example, a machine learning algorithm was developed that could infer the key based on the pitch profile of a piece. Bach's *Well-tempered Clavier* (Book 2) [23] was chosen as a test set, as it has two pieces in each of the 24 keys, providing an ordered pattern of tonality.
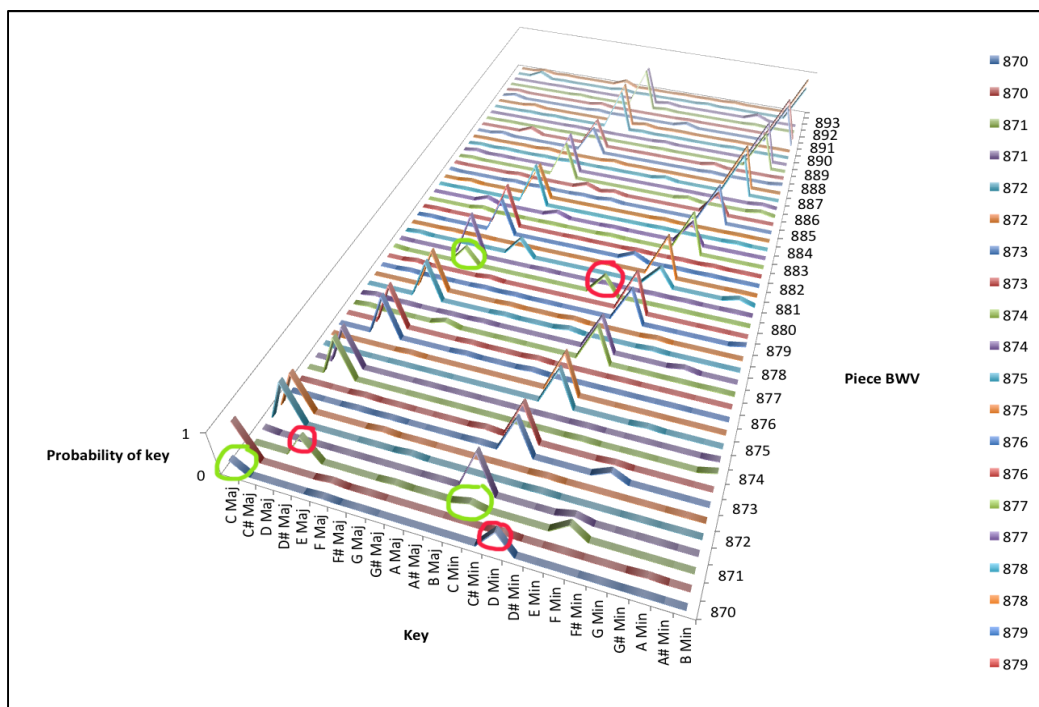
Figure 14. visualisation of a machine learning algorithms prediction of the 48 pieces of Bach's well-tempered clavier book 2 [23]. The 3 mistakes are BWV numbers 870 part 1, 871 part 1 and 880 part 1. The red highlighting shows the mistakes and the green shows the actual keys.

Figure 14 presents the detection results of the model, for each piece, ordered by their BWV number. The results of this experiment show that the algorithm is mostly able to predict each of the keys, and the graph can be inspected to find the relative confidence of each prediction – as well as identify anomalies and deviations from the expected results. Bach's methodical progression through alternating major and minor keys within the collected work produces a visual pattern in the plot (discernable from the gestalt law of good continuance [22]), the deviations from which identify errors in the key detection model – and, in turn, nuances in Bach's approach to key.

The algorithm makes three mistakes, out of a total of 48 predictions, corresponding to the anomalies circled in the figure. In one instance, the algorithm has predicted a key of D minor when the nominal key is C-Major. By visualizing the pitch profile of the piece (Figure 15), using the techniques suggested in Section 4.3 it can be observed that the overall ratio of pitch D, is higher than the tonic and 5th compared with an ideal plot (Figure 11), Indeed, this detection anomaly is attributable to Bach's actual use of D minor (and other keys) in the piece. [23] This indicates a limitation of the analysis technique, in conflating the pitch profile of an entire piece without sensitivity to modulation, but nonetheless raises an interesting musicological question of why this and not other pieces from the set fall foul of this limitation.
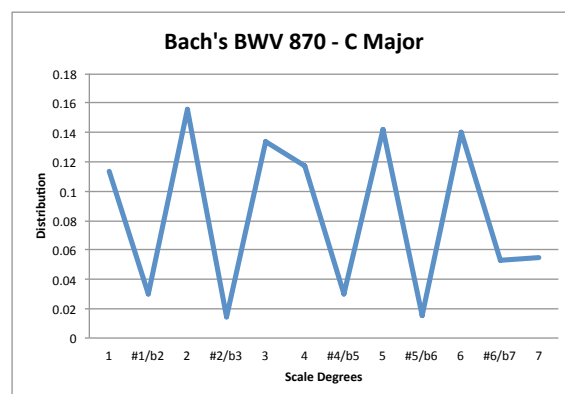

Figure 15. Pitch Distribution for Bach's BWV 870.

5. VISUALISING TIME

Visualisation can also be used to reveal patterns in musical time, as in the case of rhythm, tempo, and density. Time also provides the metrical structure to a sequence of pitches. Taking the Bach piece BWV 772 as before, and visualizing the rhythmic aspects of the piece, several patterns are revealed. The elements under consideration are Note Onset, Note Length, and Density.

5.1 Note Onset

The basic rhythmic plot, note onset (Figure 16) shows the ratio of note onsets in each position of the bar for the entire piece. The events are first quantised to 1/32nd of a note, to remove noise caused by micro variations in time.

The plot shows us, that simpler divisions of the bar are more likely to contain notes than more complex ones, shown by the regular distribution and preponderance of quavers and semi-quavers. The middle of the bar has the least note activity in general, whereas the 1st quaver beat, and 4th quaver beat have the most. Comparing this to Beethoven's piano sonata No.21 Op. 53 (Figure 17), a piece from a much later period, shows a complete contrast in the structure, with a much more uniform distribution of note onsets, with the second semi-quaver bar position (3/32) being the most likely place for a note to be played.
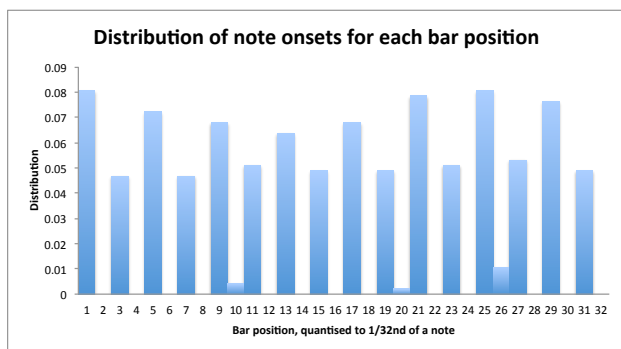

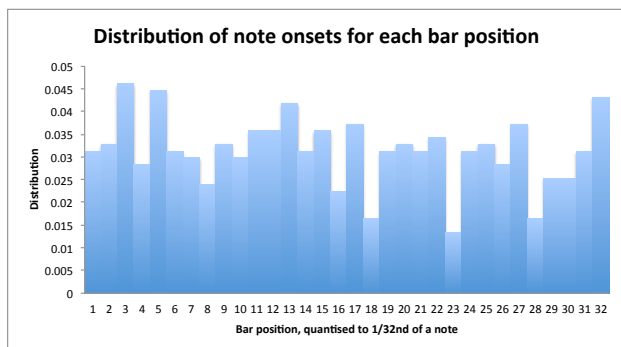
Figure 16. Distribution of rhythm for Bach's BWV 772.

Figure 17. Distribution of rhythm for Beethoven's Op. 53.

5.2 Note Length.

Note length visualisation (Figure 18) does not reveal as much information as some other techniques, but confirms this piece uses mostly note lengths of a semi-quaver in length. Roughly four times as many as using a quaver note. However, comparing this with other examples of music, for example Beethoven's piano sonata No.21 Op. 53 (Figure 19), shows for example the use of a dotted semi-quaver (3/32) note length is more common than either a quaver or crotchet, and a value not even used in the Bach piece.
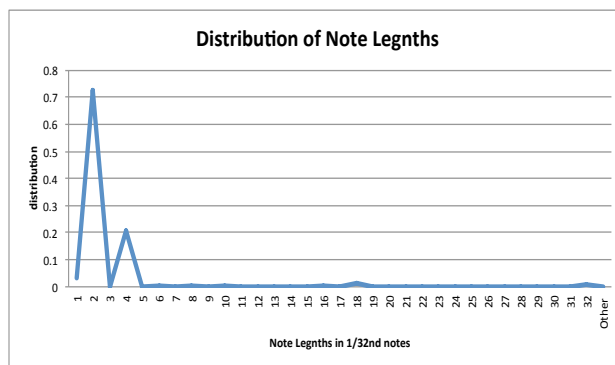

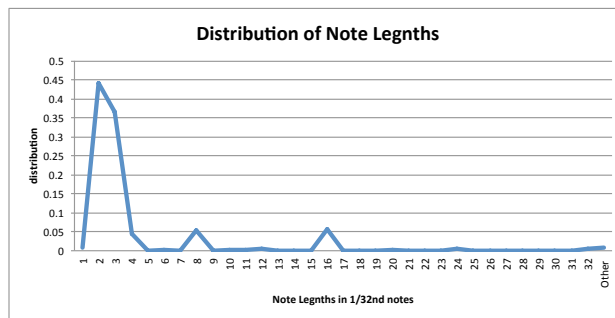**Figure 18.** Distribution of note length in Bach's BWV 772.

Figure 19. Distribution of note length in Beethoven's Op. 53.

5.3 Rhythmic Density.

Rhythmic density can be defined as the number of note onsets that happen during a beat or other window of time. The analysis is computed by calculating the number of onsets in each density window, and plotting the changes over time for each voice (note that only the first 12 measures are shown in Figure 20). Using Bach's BWV 772 again, several repeating patterns are visually observable between the two voices.

Figure 20 shows that only three of 48 windows have both voices indicating a density reading of 4 simultaneously. The sharp peak in Voice 1 at 23-24, is indicated as the most intense, a result of the piece using demi-semi-quavers (see figure 21). From windows 25 to 41, the voices are alternating in a strict pattern. This representation provides a concise overview, but does not differentiate between chords and rapid melodic phrases, reducing the *visibility* and *juxtaposability* of data. However, while a finer resolution could reveal more detail, it would also reduce *conciseness*, with four times as many data points. This represents a common trade-off between the dimensions, as observed in other notations [4].

In general, the techniques discussed in this section show that one method will reveal certain information at the sake of obscuring others, and that sometimes multiple perspectives are needed to fully understand the data.
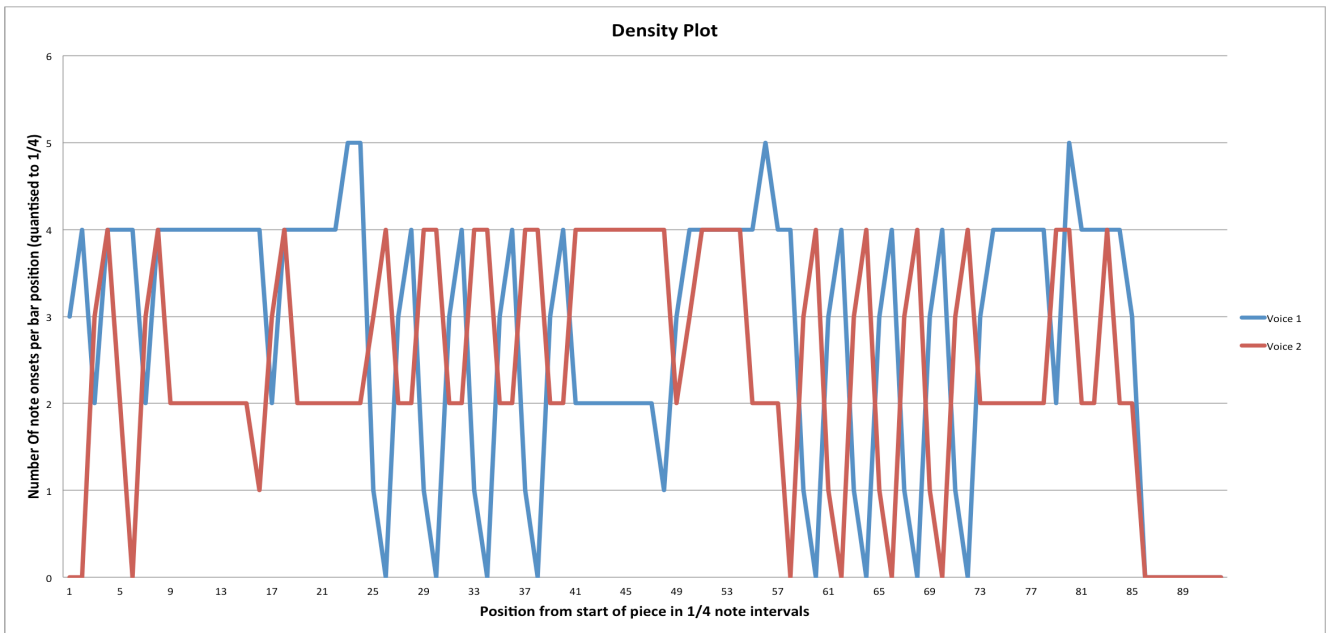
Figure 20. Density visualisation of the fist 12 bars of Bach's BWV 772, quantised to 1/4 (crotchet) note.

figure 21. Demi-semi-quavers in bars 6. Relative to points 23-24 on figure 20.

6. INTEGRATED VISUALISATIONS

Previous sections considered elements of music in isolation, but visualisations can also reveal relationships between different dimensions of music. The ability to integrate musical characteristics and model the complex interwoven principles between them is a prime objective of music analysis and visualisation. The diversity and variety of such interconnections makes this difficult, but it is possible to combine multiple dimensions of characteristics to reveal more complex and interesting patterns.

Two related elements of music that can be integrated for visualisation and analysis are rhythm and pitch. A sequence of notes can be considered a pitch change after a given length of time, and it is possible to build up the frequency of these different event combinations and display the result. Given a standard composition the number of options is vast, and represents a complex problem. However, this is relatively easy to visualise (Figure 22) by plotting the change in interval against the difference between note onset, with the colour level (brightness)

showing the ratio. In the example (Figure 22), a visualisation of Bach's Brandenburg concerto BWV 1046 [23] is shown, using this method.
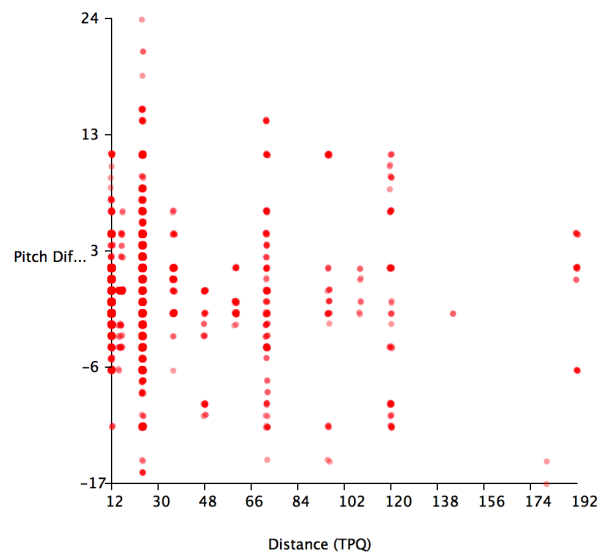

Figure 22. Visualisation of change in interval vs time between note onsets for Bach's BWV 1046. TPQ is defined as the number of ticks per quarter (crotchet) note.

Looking at the analysis, it is clear how consistent the timing of the piece is, with most events falling on quaver note divisions. There is some evidence of quaver-triplets as shown between 12 TPQ (Ticks Per Quarter Note or Crotchet) equivalent to a semi-quaver and 24 TPQ (Crochet), with these taking a value of 16 TPQ. Looking at the overall pitch range the widest range of pitch intervals is a note following on a quavers length after the previous

note, with events ranging from +24 semitones, to -17 semitones. This is also where the most events are likely to be played, shown by the density of red dots. At the 1 and 2 semi-quaver duration (12 and 24 TPQ) the pitch is more likely to increase, on any value greater than this, the pitch is likely to decrease. At the semiquaver difference, almost all intervals are present, but compare this to longer duration differences, and intervals start to disappear. An interval change of +4 semitones (major 3rd) does not happen following a previous note whose duration was a quaver. This is quite possibly linked to the rules of strict counterpoint, a technique regularly employed by the composer, but further investigation is subsequently required before drawing specific conclusions. Finally, at the 3-semi quaver duration (32 TPQ) interval, a pitch increase is more likely, but at the crotchet level (48 TPQ) a pitch decrease is more likely.

7. CONCLUSIONS

This paper has reviewed a variety of basic music visualisations to demonstrate their utility to reveal implicit details, patterns, and structures in musical phrases, pieces and broader corpora. Although the visualisations have been informally evaluated with reference to the CDMN framework, another way to evaluate the use of visualisation is to establish whether or not it revealed something that was either not known before or complicated to reveal using other methods. As several of these techniques have made such novel observations about musical structure, they can therefore be considered successful.

Other further types of studies are also planned in this area, including embedding these visualisation techniques inside music composition software. Such investigations will explore the pedagogical benefits of alternative visual representations of music, looking at how visualisations can inform students understanding of musical process and structure.

Visualisation techniques can also inform the design of generative musical techniques. They allow the identification of characteristics that can become factors of a computer composition models, such as the parameters of a machine learning process. It also allows a degree of quantitative evaluation and comparison between music generated algorithmically and the target musical result. Vickery [25] advocates re-sonifying visualised music representations, formed through analysis of the original music.

While this review of visualisation techniques only scratches the surface of both visual and musical possibilities, it is clear the visual domain can be exploited to provide different perspectives on musical patterns and structures, and make hidden information and insights more accessible to musicians and scholars.

8. REFERENCES

[1] B. Shneiderman, "The eyes have it: A task by data type taxonomy for information visualisations", in *Symposium on Visual Languages*, BouIder, Colorado, 1996, pp. 336-343.

[2] B. Craft and P. Cairns, "Beyond guidelines: what can we learn from the visual information seeking mantra?", in *Ninth International Conference on Information Visualisation*, London, UK, 2005, pp. 110--118.

[3] A. Soriano, F. Paulovich, L. Nonato and M. Oliveira, "Visualisation of music collections based on structural content similarity", in *27th SIBGRAPI Conference on Graphics, Patterns and Images*, Brazil, 2014, pp. 25-32.

[4] C. Nash, "The cognitive dimensions of music notations", in *First International Conference on Technologies for Music Notation and Representation*, Paris, France, 2016, pp. 191-203.

[5] J. Foote, "Visualizing music and audio using self-similarity", in *Proceedings of the seventh ACM international conference on Multimedia (Part 1)*, Orlando, Florida, 1999, pp. 77-80.

[6] J. Wolkowicz and S. Brooks, "Midivis: Visualizing music pieces structure via similarity matrices", in *Proceedings of the 2009 International Computer Music Conference, ICMC'09*, Montreal, Canada, 2009, pp. 53,60.

[7] T. Bergstrom, K. Karahalios and J. C. Hart, "Isochords: visualizing structure in music", in *Proceedings of Graphics Interface 2007*, Montreal, Canada, 2007, pp. 297--304.

[8] S. Holland, K. Wilkie, A. Bouwer, M. Dalgleish and P. Mulholland, "Whole body interaction in abstract domains", in *Whole body interaction*, 1st ed., D. England, Ed. London, Uk: Springer, 2016.

[9] D. Jeong and J. Nam, "Visualizing Music in its Entirety using Acoustic Features: Music Flowgram", in *Proceedings of the International Conference on Technologies for Music Notation and Representation - TENOR2016*, Cambridge Uk, 2016, pp. 25-32.

[10] D. Herremans, and E. Chew, "Tension ribbons: Quantifying and visualising tonal tension", in *Proceedings of the International Conference on Technologies for Music Notation and Representation - TENOR2016*, Cambridge Uk, 2016, pp. 8-18.

[11] F. Lerdahl and R. Jackendoff, *A generative theory of tonal music*, 1st ed. Cambridge, Mass.: MIT Press, 1983.

[12] A. Forte and S. Gilbert, *Introduction to Schenkerian analysis*, 1st ed. New York: Norton, 1982.

[13] M. Hamanaka, K. Hirata and S. Tojo, "Implementing "A Generative Theory of Tonal Music"†", *Journal of New Music Research*, vol. 35, no. 4, pp. 249-277, 2006.

[14] A. Marsden, "Schenkerian Analysis by Computer: A Proof of Concept", *Journal of New Music Research*, vol. 39, no. 3, pp. 269-289, 2010.

[15] T. R. G. Green and M. Petre, "Usability analysis of visual programming environments: A 'cognitive dimensions' framework," Journal of Visual Languages & Computing, vol. 7, no. 2, pp. 131–174, 1996.

[16] D. Temperley, *Music and probability*, 1st ed. Cambridge, Mass.: MIT Press, 2007.

[17] D. Temperley, *The cognition of basic musical structures*, 1st ed. Cambridge, Mass.: MIT Press, 2001.

[18] Y. Tomita. (1999) The Inventions and Sinfonias (1st) [online].
Available: http://www.music.qub.ac.uk/tomita/essay/inventions.html

[19] T. Fujioka, L. Trainor, B. Ross, R. Kakigi and C. Pantev, "Musical Training Enhances Automatic Encoding of Melodic Contour and Interval Structure", *Journal of Cognitive Neuroscience*, vol. 16, no. 6, pp. 1010-1021, 2004.

[20] R. Aiello and J. Sloboda, *Musical perceptions*, 1st ed. New York: Oxford University Press, 1994, pp. 173-190.

[21] M. Wood, "Visual Confusion in Piano Notation", in *Proceedings of the International Conference on Technologies for Music Notation and Representation - TENOR2016*, Cambridge Uk, 2016, pp. 230-239.

[22] G. Hartmann, *Gestalt psychology*, 1st ed. New York: The Ronald Press Company, 1935.

[23] "Dave's J.S. Bach Page - MIDI Files", *Jsbach.net*, 2016. [Online]. Available: http://www.jsbach.net/midi/. [Accessed: 21- Nov- 2016].

[24] "3 Pieces, Op.11 (Schoenberg, Arnold) - IMSLP/Petrucci Music Library: Free Public Domain Sheet Music", *Imslp.org*, 2016. [Online]. Available: https://imslp.org/wiki/3_Pieces,_Op.11_(Schoenberg,_Arnold). [Accessed: 27- Nov- 2016].

[25] L. Vickery, "Hybrid Real/Mimetic Sound Works", in *Proceedings of the International Conference on Technologies for Music Notation and Representation - TENOR2016*, Cambridge Uk, 2016, pp. 19-24.