Similarity and Categorisation in Boulez' *Parenthèse* from the Third Piano Sonata: A Formal Analysis.

Christina Anagnostopoulou ⋆ and Alan Smaill † *

⋆ Faculty of Music, University of Edinburgh

† Division of Informatics, University of Edinburgh

chrisa@music.ed.ac.uk, smaill@dai.ed.ac.uk

1 Introduction

Categorisation is the process of detecting structures and similarities between the objects in the world, and grouping similar objects together into classes. This process lies at the basis of most human cognitive activities. Equally, similarity and difference relations play a fundamental role in the internal structure of a musical piece, and in our musical understanding (e.g., [Deliège, 96]). Many theories and analytical methods in music, such as traditional morphological analysis, paradigmatic analysis, pitch class set theory, motivic analysis and so forth, are based on similarity relations.

A problem with a categorisation-based approach to music analysis is that often the categories in musical pieces are chosen intuitively, making it difficult to justify the choice of a specific class for a musical segment, and introducing inconsistencies into the analysis. In this paper we address this problem by presenting a formal approach to categorisation which is based on a clustering algorithm that operates on well-defined descriptions of musical segments, and we apply this approach to the analysis of a musical piece, namely, Boulez' *Parenthèse*, a movement from his 3rd piano sonata.

The rest of this paper is organized as follows: in the next section, we briefly describe Boulez' 3rd piano sonata and *Parenthèse*, and we discuss the challenges that this piece poses to the analyst. Then, we explain in detail our formal approach

*Many thanks to Gert Westermann and Fred Howell.

to categorisation, including segmentation of the piece, representing the segments in terms of musical features, and clustering these representations with a computational algorithm. Section 4 describes the categorisation experiments that were carried out, and the results of these experiments are presented in section 5. In section 6 we discuss these results and suggest directions of future research.

2 Boulez' 3^{rd} Piano Sonata

Pierre Boulez' 3^{rd} piano sonata is based on difference as much as on similarity. There are various strong relationships between the movements that need not concern us here, where we aim to study the low-level relationships of a single movement. According to Ivanka Stoianova [Stoianova, 1978], 'repetition is vital, although it is "a repetition-difference" within the circumstances of the serial writing [...] In reality, it is a different kind of repetition, which is the principal generator of dodecaphony and serialism.'

Parenthèse consists of 6 fragments of music that are obligatory to play, and in between them are 5 fragments of music in parentheses, which are optional to play. According to Stoianova [Stoianova, 1978, p.140], *Parenthèse* is the

> "microcosm of the symmetrical structure of the whole sonata. The presence of the obligatory and optional (in parentheses) fragments implies the co-existence of two symmetric structures: a circular symmetry of the obligatory groups and another similar one of the groups in parentheses."

In order to capture this structure, the analysis of the entire piece can be split into three parts: first, the analysis of the six obligatory fragments, second, the analysis of the optional fragments in parentheses, and third, the relation between obligatory and optional fragments. In this paper, we demonstrate a full analysis of the first part, that is, the obligatory fragments of the piece.

Within *Parenthèse* we observe different similarity relations between its segments: first, the dodecaphonic "repetition-difference", which is based on the use of pitch class sets, and second, similarity relations in musical properties such as rhythm and tempo, tonal centres, intervals, contour, and way of playing.

The method of analysis that we present in this paper aims to bring out these relations. The aim is, on the one hand, piece-specific: to demonstrate the structure of the obligatory part of the piece. On the other hand, a more general aim is to demonstrate how the formal method of analysis, that has previously been shown to work for monophonic pieces ([Anagnostopoulou and Westermann, 97]) can be applied to a non-monophonic, atonal piece of music with very rich internal relations, and where a hierarchical segmentation is needed.

3 The Analysis

The analysis method is a formalised and extended version of Paradigmatic Analysis [Ruwet, 1996; Nattiez, 1975]. The formalisation consists in dividing the analysis process into discrete steps, fully specifiying the representations at each step, and performing the clustering of the musical segments with a well-defined algorithm.

The analysis process is illustrated in figure 1. First the piece[1] is broken down hierarchically into smaller segments, and then each of the segments is described as a set of properties. The description of the segments is then turned into an appropriate computational input in the form of feature vectors, and the classification algorithm takes this input and produces a hierarchic classification of the segments. The result of this process is a categorisation analysis that makes similarity relations explicit. In the following sections, we describe each step in detail.

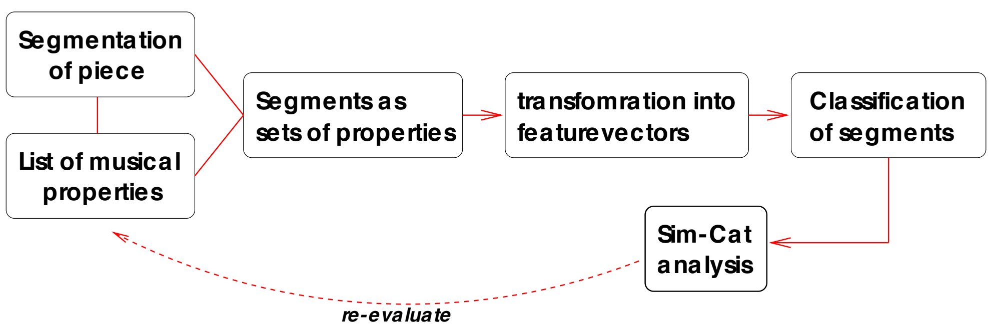

Figure 1: A general overview of the Similarity and Categorisation Method of Analysis.

3.1 Segmentation

In most formal methods of analysis, the music piece is first split into segments. It is important to consider that the precise way in which the piece is segmented has a profound influence on the outcome of the analysis.

[1]By using the term *piece* we mean the obligatory fragments that are analysed here.

In *Parenthèse*, segmentation is an easier task than for most pieces, since in most places the segmentation points are clearly indicated by the composer. We define segment boundaries

- at the beginning and end of the fragments in parentheses
- where the piano stave is marked with *V* as a break point
- where there is a more or less obvious change of texture, that is, between segments 2a and 2b, and 4c and 4d. This also coincides with the change of a pitch-class set, and this segmentation therefore corresponds to the so-called imbrication method of segmentation [Forte, 73].

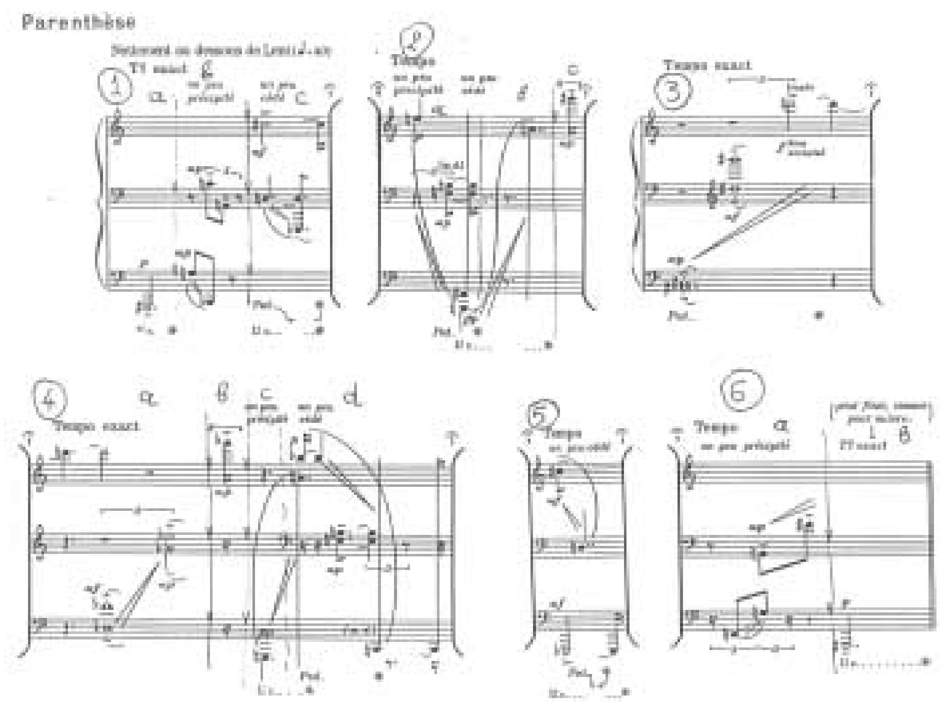

Figure 2: The obligatory fragments of *Parenthèse* and their sub-segments.

The resulting segmentation of the piece is shown in figure 2. We denote the obligatory fragments with numbers 1,. . . , 6. These fragments are then further divided into segments 1a, 1b, 1c, 2a, 2b, 2c and so on.

In the following experiments, we use three levels of segments: the undivided high-level segments 1,. . . , 6, the low level segments 1a, 1b, and so on, and an intermediate level where we combine certain adjacent low-level segments: for example, the low level segments 1a and 1b form the intermediate level segment 1ab. By this we hope to capture similarities that exist between the different segmentation levels.

3.2 Description of Segments as Sets of Properties

The term *property* is often used interchangeably with *feature* and *attribute*, and here we use it in the same way. A property value can denote the presence or absence of a feature in a segment (e.g., crescendo), or it denotes an attribute that can take one of several (mutually exclusive) values, e.g., key. In order to translate such a multi-valued property into a binary form which is required for the classification algorithm, we use a 1-out-ot-n encoding: out of the n possible values, the one that is present is set to 1, and all others to 0.

Like the choice of segmentation, the choice of properties to describe the segments has a profound influence on the results of the computational classification: the algorithm groups the segments according to their similarity, and this similarity is determined by the property values for each segment. What makes two music segments identical, similar or different will obviously depend crucially on the property selection, and on the way of representing the properties. Whereas the choice of properties is made by the analyst, the formal method of analysis shows precisely how this choice influences the resulting analysis.

In developing a set of properties, a segment is analysed in terms of various musical properties that seem important for its description and for its differentiation to other segments. Then, all properties that have been chosen for the segments are combined into one set, and each segment is described in terms of this set of properties.

The description of a segment by a list of properties is not complete: it is impossible, based on the properties, to re-create the particular segment they describe. Instead, the proeprties contain all information about a segment that are considered important for the further analysis of the piece. Different analyses warrant different properties: for example, a rhythmic analysis would describe the rhythmic properties of each segment in detail, and an analysis aiming to compare certain features across a wider music repertoire would emphasize those features.

Two kinds of properties can be used for describing a musical segment:

- properties that are true for a part of the segment, for example, the existence of a specific interval in the segment

- properties that are true for the whole of the segment, for example a rising melodic movement

In our approach we mainly make use of the second kind of properties, with the exception of specific rhythmic and intervallic patterns that describe merely part of a segment.

Table 1 shows the properties that we use in the analysis, and the segments in which they are found. The properties considered here are:

- the existence of various pitch-class sets and certain common subsets that they share. The reason to consider the common subsets is to reinforce similarity between the sets that the composer uses, which are indeed very similar to each other. In order for a pitch-class set to be true for a segment, *all* the notes of the segment have to belong to the pitch class set.
- the existence of various rhythmic patterns. These in contrast do not require for all the notes of the segment to belong to the specific rhythmic pattern.
- tempo and dynamic descriptions. The composer is very specific about which tempo and dynamic descriptions he uses, and these are important for the distinction of the segments and the overall structure of the piece, so in a classification task of this piece, they should be taken into account.
- tonal centres, which in this case are single tones rather than keys, and relations between tones, significant intervals that the composer seems to favour, and contour information.

Table 1 shows how each segment is “translated” from musical notation to a set of properties. The reason for this transformation is to achieve, at a next step, a consistent classification. For this reason we need to have the criteria set forth before the classification takes place.

Describing the segments in terms of properties (*cf.* table 1) results in a 34-bit feature vector for each segment, making it thus appropriate computational input for the classification module.

3.3 Classification

The classification of the segments that are represented as feature vectors, is carried out with a computational algorithm. This approach differs from the traditional Ruwet/Nattiez Paradigmatic Analysis in that

property	1a	1b	1c	2a	2b	2c	3	4a	4b	4c	4d	5	6a	6b
3-1(12)			y			y						y		
4-1(12)	y-	-y-	-y											
7-2				y-	-y-	-y			y-	-y-	-y			
6-9				y-	-y			y-	-y	y-	-y			
5-2	y-	-y					y	y					y-	-y
5-5				y							y			
all								y-	-y-	-y-	-y			
inv 012	y-	-y	y	y	-y-	y	y	y	-y	y-	y	y	y-	-y
inv +3	y-	-y	-y	y	-y-	-y	y	y	-y	y-	y		y-	-y
inv +5	y-	-y		y-	-y-	-y	y	y	-y	y-	-y		y-	-y
inv +7				y	-y-	-y		y-	-y	y-	y			
longn	y		y				y	y				y		y
Q,Q		y											y	
4note		y											y	
SQdot						y			y					
triplex		y		y							y		y	
exact	y						y	y	y					y
précip		y		y						y			y	
cédé			y		y	y					y	y		
mf+			y				y	y				y		
cresc	y-	-y-	-y				y					y	y	
dimin			y					y					y-	-y
steady	y	y		y?-	-y-	-y			y	y	y			
Gis/Aes	y					y			y					y
G,Gis,A	y-	-y			y-	-y			y-	-y-	-y		y-	-y
D							y	y				y		
Cis,D,Dis			y				y	y				y		
semit		y	y	y	y-	-y	y	y			y	y	y	
tritone		y		y			y				y			
third		y		y			y	y					y	
wob		y											y	
down1			y					y				y		
down2				y							y			
up2							y							

Table 1: The lowest-level obligatory segments (1a, ..., 6b) and the properties that are true for each segment. When a property exists in a segment, then this is marked by a “y”. When there is a property that is true for a bigger segment but not for the low-est level, then this is marked in the lowest-level segments that the bigger segment is made from, by using “y-”, “-y”, “-y-”, according to which adjecent the property is shared with. The first part of the table contains the pitch-class sets and their common subsets, the second part contains the rhythmic patterns, the third part contains the directions by the composer on tempo and way of playing, the fourth part contains tonal centres and specific intervals and the last part contains contour information.

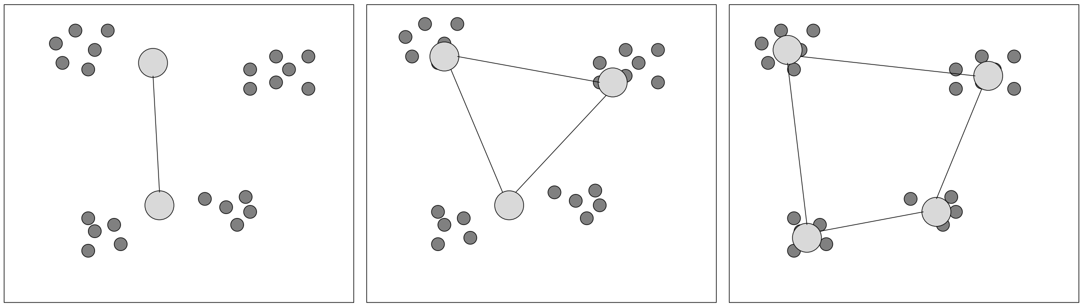

Figure 3: Construction of a GNG network. Small circles represent input data, and large circles connected with edges are the units of the network.

- the classification process is formalised and depends on explicit criteria, that is, the set of properties, thus avoiding intuitional and unfounded decisions.
- the classification proceeds in an approximately hierarchic way, from the whole piece being considered as one class to each segment being considered as a separate class.
- the algorithm develops probabilistic prototypical values of the class properties, directly showing similarities and differences between the classes.

The algorithm used here is an unsupervised neural network clustering algorithm, *Growing Neural Gas* (GNG) [Fritzke, 1995], which has been used previously for musical analysis of different musical styles and has been shown to produce valid results ([Höthker, Hörnel and Anagnostopoulou, 2000; Anagnostopoulou and Westermann, 97]).

The GNG consists of units that move towards the center of the classes, and during the learning process it adds units at a constant interval, effectively increasing the number of classes in the analysis. Each segment belongs to the class of its closest unit.

Figure 3 shows the development of a network in a two-dimensional input space with four distinct clusters. The network starts with two units and can therefore distinguish only between the two main clusters. After a certain number of presentations of the feature vectors (here 500 presentations of each vector), a new unit is inserted and the units move to the positions indicated in the second picture. When the fourth unit is inserted, the input units distribute over the four clusters.

In principal, insertion of units proceeds forever. The GNG algorithm thus lets the analyst define the level of grainedness of her analysis and does not impose *a priori* constraints on the number of clusters. Each unit forms a prototype of a cluster, expressed in the probability distribution of the feature values of their cluster members, and the distances between the units can be measured to gain information about the similarity between the classes.

4 Experiments

We performed four experiments:

In the first experiment, the classification algorithm was trained on the feature vectors that represent the segments on the smallest level only: 1a, 1b, 1c, 2a, 2b, 2c, 3, 4a, 4b, 4c, 4d, 5, 6a, 6b. The properties that stretch over adjacent smallest-level segments were not taken into account.

In the second experiment, the algorithm was again trained on feature vectors representing the smallest-level segments, but this time they were enhanced with those features that stretch over segment boundaries. For example, if segment 1 has a property *a* that is not reflected in its sub-segments 1a, 1b, and 1c, then here these sub-segments inherited this global feature.

In the third experiment, all segmentation levels were represented in parallel and the algorithm was trained on the full set of lowest-level segments 1a,. . . ,6b, the highest level segments 1,. . . ,6, and certain middle-level segments such as 1ab, 4bcd, and so on. In contrast to experiment 2, the lowest-level segments were only represented by their own properties and not the shared ones.

In the fourth experiment, we considered only a selection of eight segments drawn from all the levels: 1ab, 1c, 2, 3, 4a, 4bcd, 5 and 6.

By comparing the developing network architecture over a period of insertion of units, we were able to observe the hierarchy of classes.

5 Results

Table 2 shows the results of experiments 1 and 2, when the number of classes is 5 (that is, when the network has inserted 5 units). Table 3 shows the results in the same two experiments, when there are 7 and 8 classes.

The results of experiments 1 and 2 are all intuitively acceptable, although those from experiment 2 seem slightly better. Table 2 shows the results of experiment 1 with 5 classes: here, 1a, 2b, 2c, 4b, 4c, 6b belong to the same class. This classification would have been better if segments 1a and 6b were in a different class from the others, since they contain long notes whereas the other segments contain shorter notes. This difference could be enhanced by introducing an extra

Class	Exp 1	Exp 2
Class I	2a, 4d	2a, 2b, 2c, 4b, 4c, 4d
Class II	3, 4a	1c, 5a
Class III	1a, 2b, 2c, 4b, 4c, 6b	1b, 6a
Class IV	1b, 6a	3, 4a
Class V	1c, 5a	1a, 6b

Table 2: The experimental results in the two first experiemtns when the number of classes is 5.

Class	Exp 1	Exp 1	Exp 2
Class I	2a, 4d	2a, 4d	2c, 4b
Class II	3, 4a	3, 4a	1c, 5a
Class III	1a, 6b, 4b	1a, 6b	1b, 6a
Class IV	1b, 6a	1b, 6a	1a, 6b
Class V	1c, 5a	1c, 5a	3, 4a
Class VI	2b, 4c	2b, 4c	2a, 4d
Class VII	2c	2c	2b, 4c
Class VIII	–	4b	–

Table 3: The experimental results in the 2 experiments when the number of classes is 7 or 8.

feature *note-length* in the list of properties describing the segments. This is an example of re-evaluation of the segment descriptions.

Table 3 shows the classification for experiment 1 with 7 and 8 classes. Here we see that the same segments are separated into three classes when the overall number of classes is 7. Therefore a bigger number of classes produces more satisfactory results in this case.

Whereas experiment 1, which does not incorporate properties that stretch over segment boundaries, emphasised the iconic similarity between segments, in experiment 2 the structural similarity between segments is enhanced due to the added "global" features relating to higher-level segments. Here, all the subsegments of segments 2 and 4 are in the same class. Even though the iconic similarity of these segments is low (e.g., 4b and 4d), they both share the global properties of segment 4.

Figure 4 shows the progression of the classification in experiment 2 from 2 to 10 classes. This is an intuitively successful example of hierarchic classification, which shows the symmetrical structure of the piece.

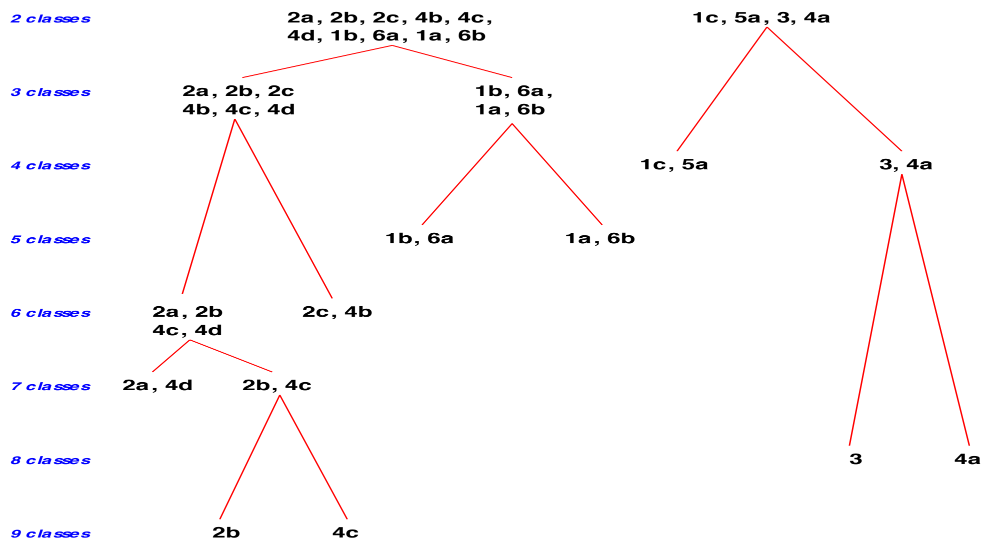

Figure 4: Hierarchic classification for set 2.

In experiment 3 all levels of segments are taken into account. The results for 5 and 8 classes are shown in table 4. In this case we often get segments and their subsegments classified in the same category, since they share many of their properties (for example, segments and subsegments of 2 and 4). This problem cannot be avoided in such a setting and the results need further interpretation in order to be valid. For this reason, 5 classes seems to be too few classes for an acceptable classification. When the number of classes increases to 8, the results improve: 3 and 4a are correctly classified into a category of their own, and the same holds for 1b and 6a. It is interesting to see segment 4 on a category of its own, since it is the longest segment of all. Segments 2 and 4bcd are placed in the same category and are an example of similarity across levels. In general, 8 classes seem to be sufficient for demonstrating the symmetry of the segments, although one needs to consider carefully which segments denote this and which are merely related subsegments of the same bigger segment.

Experiment 4 is the simplest experiment because we consider only a selection of 8 segments across levels. These are chosen in order to show the structure of the piece. Table 5 shows the resulting classification when having 4 classes: the first and last segment, 1ab and 6, are classified together, and the same is true for 1c and 5, 2 and 4bcd and 3 and 4a. These segments are almost mirror images of each other, and define the symmetrical structure of the piece.

Class	Exp 3, classes 5	Exp 3, classes 8
Class I	1c, 3, 4a, 4, 4ab	1c, 5a, 4ab
Class II	1a, 2b, 2c, 4b, 4c, 6b, 2bc	1a, 2b, 2c, 4b, 4c, 6b
Class III	2a, 4d, 2, 2ab, 4cd, 4bcd	1, 6, 1ab, 1bc
Class IV	1b, 6a, 1, 6, 1ab	2a, 4d, 2, 2ab, 4cd, 4bcd
Class V	1bc	4
Class VI		1b, 6a
Class VII		2bc
Class VIII		3, 4a

Table 4: The experimental results in the third experiment when the number of classes is 5 and 8.

Class	Experiment 4
Class I	1c, 5
Class II	2, 4bcd
Class III	1ab, 6
Class IV	3, 4a

Table 5: The experimental results in experiment 4, with 4 classes.

6 Conclusions

We presented a formal method of analysis based on categorisation of music segments according to similarity. We applied this method to the analysis of Boulez' *Parenthèse* from the 3rd piano sonata, taking into account the obligatory fragments of the piece. The resulting hierarchic classification defines the similarity and difference relations between classes and between segments. We demonstrated how a classification analysis is appropriate for this piece and how it brings out the symmetrical structure that the composer intended. This method of analysis, shown previously to work on more traditional kinds of music, is shown here to be appropriate for an atonal and non-monophonic piece of music.

The results give many interesting insights on the obligatory fragments. In terms of internal relations, it is a very rich piece, each note situated in its position for a variety of reasons, forming part of an overall plan. More specifically, we see that the piece also has an interesting tonal structure, evolving mainly around G sharp at the beginning and end, and around D in the middle of the piece. The pitch class sets used are very similar to each other, segments 2 and 4 sharing sets, and the same for segments 1, 3 and 6. Dynamics and tempo seem to be very important for the segmentation and difference between subsegments, whereas contour information seems to be reflecting the symmetrical structure of the piece.

The issue of hierarchic segmentation in a classification task poses interesting challenges to the analyst. When classifying all the levels at the same time, on the one hand we get interesting similarities across levels, but on the other hand we get similarities between segments and their subsegments which are redundant.

The results depend on the initial representation, that is the choice of properties according on which each segment is described. A different choice of properties would yield different results. However, a bad resulting classification would show that the initial properties were not chosen carefully, and a re-evaluation of these properties is needed. In that way, the analyst can revise the initial properties. This procedure can go on until an acceptable classification is produced.

The principles of similarity and difference are common principles to the vast majority of musical repertoire. It can be argued that they are responsible to a large extent for cohesion and coherence within the musical piece.

7 Bibliography

Anagnostopoulou, C. and Westermann, G, 1997, *Classification in Music: A Formal Model for Paradigmatic Analysis*, Proceedings of the International Computer Music Conference, Thessaloniki, Greece.

Deliège, I., 1996, Cue Abstraction as a Component of Categorisation Processes in Music Listening, in Psychology of Music, 20(2).

Forte, A., 1973, *The Structure of Atonal Music*, Yale University Press, New Haven and London.

Fritzke, B., 1995, *A Growing Neural Gas Algorithm Learns Topologies*, in G. Tesauro and D. S. Touretzky and T. K. Lean (eds), Advances in Neural Information Processing Systems 7, MIT Press.

Höthker, K., Hörnel, D. and Anagnostopoulou, C., 2000, *Investigating the Infuence of Representations and Algorithms in Music Classification*, in Computers and the Humanities, 34(4), in press.

Nattiez, J.J., 1975, *Fondements d'une Sémiologie de la Musique*, Union Générale d'Editions, Paris.

Stoianova, I., 1978, *Geste-Texte-Musique*, Union Générale d'Editions, Paris.

Ruwet, N., 1966, *Méthodes D'Analyse en Musicologie* in Revue Belge de Musicologie 20.