Discovery of Musical Patterns and Their Variations with Mathematical Morphology

NILS QUAETAERT

Master in Mathematics
School of Engineering Sciences
Date: March 14, 2025
Supervisors: Paul Lascabettes, Isabelle Bloch, Moreno Andreatta and Kathlén Kohn
Examiner: Kathlén Kohn
Host organization: IRMA, Strasbourg, France
Swedish title: Metod för att Upptäcka Musikaliska Mönster och Deras Variationer med Matematisk Morfologi

Abstract

Identifying the musical motifs present in a composition is a fundamental aspect of our current understanding of musical analysis. Using mathematical formalism to discover musically significant recurring patterns in a piece can thus help to assist and refine musical analysis. While first attempts to deal with this problem have been based on a string-type representation of music, further developments have shown the advantages of using a geometric representation. Using such a representation, point-set algorithms can be designed in order to discover certain classes of patterns within a piece. Furthermore, recent developments have shown that the tools of mathematical morphology can be used to devise such algorithms. So far, mathematical morphology has been used for the discovery of patterns and their exact occurrences within a composition. In this master's thesis, we propose a mathematical formalism with the purpose of discovering patterns and their occurrences with variations in pitch and rhythm. We define three operations based on morphological erosion: the Opening-less Variational Erosion (or OVE), the Pivotal Variational Erosion (or PVE) and the Intersectional Variational Erosion (or IVE). Each of these operations holds properties that allow us to find patterns and their occurrences with variations in a composition. We then apply our findings to a specific case: the fugues from Bach's Well-Tempered Clavier. Two algorithms, one based on the OVE and the other based on the PVE, are designed in order to identify the subject of the fugues. Both algorithms are able to do so consistently, as well as identifying some relevant truncated instances of the subject.

Keywords: Musical Motifs, Musical Variations, Mathematical Morphology, Point-Set algorithms, The Well-Tempered Clavier.

Sammanfattning

Att identifiera de musikaliska motiv som finns i en komposition är en grundläggande aspekt av vår nuvarande förståelse av musikalisk analys. Att använda matematisk formalism för att upptäcka betydelsefulla återkommande mönster i ett musikverk kan därför bidra till att underlätta och förfina musikalisk analys. De första försöken att hantera detta problem har baserats på en strängbaserad representation av musik, men vidareutveckling har visat fördelarna med att använda en geometrisk representation. Med hjälp av en sådan representation kan algoritmer utformas för att upptäcka vissa klasser av mönster i ett verk. Den senaste utvecklingen har dessutom visat att verktygen inom matematisk morfologi kan användas för att ta fram sådana algoritmer. Hittills har matematisk morfologi använts för att upptäcka mönster och deras exakta förekomst i en komposition. I denna examensarbete föreslår vi en matematisk formalism i syfte att upptäcka mönster och deras förekomst med variationer i tonhöjd och rytm. Vi definierar tre operationer baserade på morfologisk erosion: *Openingless Variational Erosion* (eller OVE), *Pivotal Variational Erosion* (eller PVE) och *Intersectional Variational Erosion* (eller IVE). Var och en av dessa operationer har egenskaper som gör att vi kan hitta mönster och deras förekomst med variationer i en komposition. Vi tillämpar sedan våra resultat på ett specifikt fall: fugorna från Bachs *Wohltemperirte Clavier*. Två algoritmer, den ena baserad på OVE och den andra baserad på PVE, är utformade för att identifiera subjektet i fugorna. Båda dessa algoritmer kan göra det konsekvent och även identifiera några relevanta trunkerade instanser av subjektet.

Nyckelord: Musikaliska motiv, Musikaliska variationer, Matematisk morfologi, Punktuppsättningsalgoritmer, Das Wohltemperierte Klavier.

Contents

Chapter 1

Introduction: Musical Patterns and Mathematical Spaces

In the first chapter of his seminal book *Harmony* [1], H. Schenker identifies the motif as the fundamental building block of music as an art form. For the music theorist, the underlying idea behind the motif is that of repetition: "only by repetition can a series of tones be characterized as something definite" (p.5). It is only through repetition that a bundle of notes can become identifiable as an individual motif, and be endowed with some sort of associative and aesthetic quality.

Though the work of H. Schenker has been undergoing critical reevaluations in academic circles for its overreliance on the western classical canon to explain music as a whole, it seems that this fundamental idea as become ingrained in our current understanding of music, as many music psychologists and music analysts have underlined that identifying the significant repetitions in a piece of music is essential to achieve a meaningful interpretation of it [2, 3, 4, 5, 6].

Therefore, the task of finding repeating patterns in a musical piece from its symbolic representation is a meaningful one. By *symbolic representation*, we designate any set of symbols that communicates the way a musical piece should be performed, e.g. sheet music or tablature. Knowing how to perform such a task could be used to assist musical analysis and detect previously unseen underlying structures in some pieces. If that is what we see to accomplish, the first question we need to ask ourselves is of course: what is a repetition ? More precisely, when can we consider that two sets of tones are perceived as "the same" ?

Such a broad question, which would mobilize notions in psychology, acoustics and music theory is evidently outside of the scope of this thesis. However, for our purposes, we point out that a repetition can be generally viewed as a *transformation*. The most straightforward type of such a transformation is of course the transposition, where a set of tones is transcribed in a higher or lower register. But musical motifs can also be truncated, augmented, diminished, inverted, reversed, embellished, and so on. As such, we can already identify a link between the notion of repetition in music and geometry through this broad concept of "transformation". In this thesis, we are mainly interested in tonal and rhythmic variations of a musical pattern, namely the case of repetition where of the pattern (or some of its notes) presents some variations either in tone or in rhythm.

The problem of finding the repetition of musical patterns is not a new one. The first attempts at solving this problem rest upon a string-based approach, wherein it is assumed that the music to be analyzed is represented by a string of symbols. The algorithms developed within this framework generally suffer from a few weaknesses. First, their scope are generally limited to monophonic sources and voiced polyphonic sources (a.k.a music that can be viewed as a superposition of voices), and they cannot be used for unvoiced polyphonic music. Thus, though they can be very useful for some specific types of music, such as fugues for example, they cannot be used in a general case. Second, to compute the similarity between patterns, they generally use the *edit-distance* which is typically not capable of finding a match between a pattern and a highly-embellished variation of it. A typical example of such an approach can be found in P.Y. Rolland's FlExPat program [7].

The introduction of a multi-dimensional framework, wherein every note is represented by a point in a space, allowed for better results. By defining classes of repeating patterns, such as Maximal Translatable Patterns (MTP) or Maximal Translational Equivalence Classes (MTEC), one can devise methods in order to find elements of those classes. Notably, many such methods are based on the SIA algorithm, which is able to compute the MTPs of a point set [8]. This approach is generally combined with a set of heuristics that allows us to select musically interesting repeating patterns.

Recent developments, notably the work of P. Lascabettes, have brought to light the link between the point-set algorithms and mathematical morphology, a theory based on geometry and lattice theory generally employed for analyzing geometrical structures in images. More specifically, it was shown that patterns could be found using a morphological operation called *erosion*. What is particularly interesting with this approach, is that it highlighted the duality between the problems of finding a pattern and its occurrences. Morphological erosion gives us an operation that loops between the set of points describing a pattern and the set of points describing its occurrences, meaning that you can deduce one from the other.

It is within this framework of mathematical morphology that our contribution falls. Indeed, one limitation we meet with this approach is that it doesn't take into account variations in the pattern. Morphological erosion allows us to find transpositions of a pattern within a musical piece, but as we have pointed out before, the repetition of a pattern encompasses other types of transformations. Our main research question for this thesis is therefore the following: how can we find occurrences of a pattern with variations using the tools of mathematical morphology ? The objective is to provide a framework that is able find musical motifs and their variations given a certain approximation while keeping the looping structure that allows to deduce a pattern from its occurrences and vice-versa.

Chapter 2

State of the Art

This chapter reviews the main topics that constitute the framework for our work. We describe the basic operations of mathematical morphology, point-set algorithms for the discovery of musical patterns, and the link between those two fields.

2.1 Mathematical Morphology: a Theory for Analyzing Geometrical Structures

In *Handbook of Spatial Logics* [9], it is said that mathematical morphology arose in 1964 as a branch of image processing, with its main concepts and tools developed in the works of G. Matheron and J. Serra [10, 11, 12]. It borrows concepts from various branches of mathematics, such as algebra, topology and discrete geometry, with the goal of analyzing the shape and topology of digital objects.

The theory we are mainly focusing on is "binary" mathematical morphology. In this case, we define a set E where E is of the form $\mathbb{R}^n$, $\mathbb{Z}^n$, or $\mathbb{Z}_n$. The power set $\mathcal{P}(E)$, provided with the inclusion relation which endows it with a lattice structure, is the main point of focus.

To begin with morphology, it is useful to fix the following notations:

- The complement of X is $X^c = \{x \in E \mid x \notin X\} = E \setminus X$.
- The translation of X by $t \in E$ is $X_t = \{x + t \mid x \in X\}$.
- The symmetrical of X by $t \in E$ is $\check{X} = \{-x \mid x \in X\}$ (when E =$\mathbb{R}^n$ or E = $\mathbb{Z}^n$).

2.1.1 Dilation and Erosion

Binary mathematical morphology is based on operations that allow for the transformation and filtering of sets of points. Basic morphological operations on sets can be obtained by combining set-theoretical operations with two basic operators, dilation and erosion, which are based on the Minkowski additions and subtractions.

Definition 2.1.1. *Let $P \in \mathcal{P}(E)$, the dilation δ_P and erosion ε_P by P are defined as:*

$$\begin{array}{rcccccl} \delta_P & : & \mathcal{P}(E) & \to & \mathcal{P}(E) & & \\ & & X & \mapsto & X \oplus P & = & \{x + p \mid x \in X, p \in P\} \\ & & & & & = & \{x \in E \mid \check{P}_x \cap X \neq \emptyset\} \end{array}$$

$$\begin{array}{rcccccl} \varepsilon_P & : & \mathcal{P}(E) & \to & \mathcal{P}(E) & & \\ & & X & \mapsto & X \ominus P & = & \{x \in E \mid P_x \subseteq X\} \end{array}$$

In this definition, though X and P play, formally speaking, similar roles, in practice they generally correspond to objects with different characteristics. X is thus oftentimes called the *Image* (a set that is generally big and given by the problem) and P the *Structuring Element* (a set that is generally small and chosen by the problem-solver). The set X^c is also given a name: it is called the *Background*.

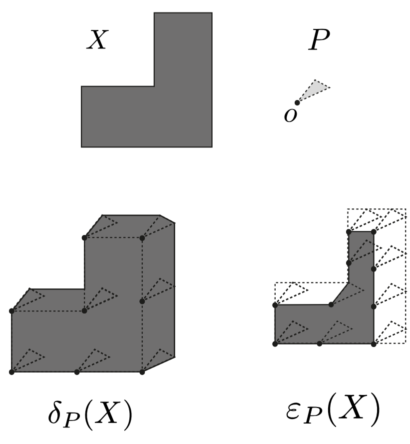

Figure 2.1: Top: The image X and the structuring element P. Bottom: the dilation $\delta_P(X)$ of X by P and the erosion $\varepsilon_P(X)$ of X by P.

It should also be noted that the structuring element is defined in relation to the origin O_E of the space. This has a direct impact on the Minkowski additions and subtractions, so when choosing a structuring element, one should not only look at its shape but also its placement in the space.

These operations can be interpreted in the following way:

- the dilation of X by P corresponds to the union of all translations of P from positions in X;
- the erosion of X by P corresponds to all the positions at which P occurs in X.

These two operations are dual by complementarity: dilating a set is equivalent to eroding its complement with the symmetric of the structuring element. In other words, eroding an object is akin to dilating the background.

Proposition 2.1.1. *Dilation and erosion are dual by complementarity [9]:*

$$(\delta_P(X))^c = \varepsilon_{\check{P}}(X^c) \text{ and } (\varepsilon_P(X))^c = \delta_{\check{P}}(X^c)$$

.

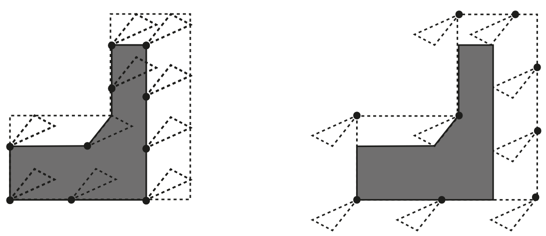

Figure 2.2: An illustration of duality between dilation and erosion. The erosion of the image (left) is equivalent to the dilation of the background by the symmetric of the structuring element (right).

With this duality in mind, we can see that both operations are intrinsically linked: in fact, the properties of erosion can be derived from the properties of dilation through duality. For instance, since dilation inflates the object and deforms convex corners of the object, we can deduce that erosion deflates the object and deforms concave corners of the object.

Beyond this, these basic morphological operations present several mathematical properties that must be pointed out, as they are generally useful for proofs and understanding their behavior as filters.

- Concerning extensiveness, both operations present properties provided that the origin is part of the structuring element:

  The dilation by P is extensive ($X \subset \delta_P(X)$) if and only if $0_E \in P$
  The erosion by P is anti-extensive ($\varepsilon_P(X) \subset X$ if and only if $0_E \in P$)

- Dilation and erosion are increasing:

$$X \subset Y \Rightarrow \delta_P(X) \subset \delta_P(Y)$$
$$X \subset Y \Rightarrow \varepsilon_P(X) \subset \varepsilon_P(Y)$$

- Dilation is increasing according to the structuring element and erosion is decreasing according to the structuring element:

$$P_1 \subset P_2 \Rightarrow \delta_{P_1}(X) \subset \delta_{P_2}(X)$$
$$P_1 \subset P_2 \Rightarrow \varepsilon_{P_2}(X) \subset \varepsilon_{P_1}(x)$$

- Concerning the union and intersection of sets:

$$\delta_P(X \cup Y) = \delta_P(X) \cup \delta_P(Y) \quad \text{and} \quad \delta_P(X \cap Y) \subset \delta_P(X) \cap \delta_P(Y)$$
$$\varepsilon_P(X \cap Y) = \varepsilon_P(X) \cap \varepsilon_P(Y) \quad \text{and} \quad \varepsilon_P(X \cup Y) \supset \varepsilon_P(X) \cup \delta_P(Y)$$

- The dilation, when seen as an operator with two arguments, is commutative with respect to the image and the structuring element:

$$\delta_P(X) = \delta_X(P)$$

- These two operations verify the *adjunction property*, which is significant in a lattice-theoretical framework:

$$\delta_P(X) \subseteq X' \iff X \subseteq \varepsilon_P(X')$$

- These operations verify the *iteration property*:

$$\delta_P(\delta_{P'}(X)) = \delta_{P \oplus P'}(X) \quad \text{and} \quad \varepsilon_P(\varepsilon_{P'}(X)) = \varepsilon_{P \oplus P'}(X)$$

2.1.2 Opening and Closing

Along with dilation and erosion, we can identify two other fundamental operators that can be derived from the previous two: opening and closing. They are defined as follows:

Definition 2.1.2. *Let $P \in \mathcal{P}(E)$. The opening γ_P and closing φ_P by P are defined by:*

$$
\begin{array}{rcccccl}
\gamma_P & : & \mathcal{P}(E) & \to & \mathcal{P}(E) & & \\
 & & X & \mapsto & X \circ P & = & (X \ominus P) \oplus P
\end{array}
$$

$$
\begin{array}{rcccccl}
\varphi_P & : & \mathcal{P}(E) & \to & \mathcal{P}(E) & & \\
 & & X & \mapsto & X \bullet P & = & (X \oplus P) \ominus P
\end{array}
$$

These operations can be interpreted in the following way:

- the opening of X by P corresponds to the union of all occurrences of translations of P in X;
- the closing of X by P corresponds to all positions at which P occurs in the union of all translations of P from positions in X.

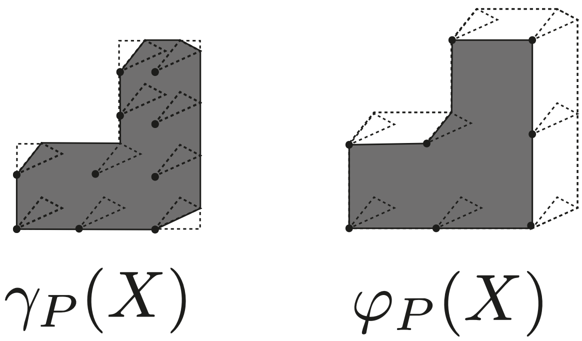

Figure 2.3: Left: Opening of X by P. Right: Closing of X by P.

Proposition 2.1.2. *Opening and closing are dual by complementarity [9]:*

$$(\gamma_P(X))^c = \varphi_{\check{P}}(X^c) \ \textit{and} \ (\varphi_P(X))^c = \gamma_{\check{P}}(X^c)$$

Hence, the properties of closing are derived from the properties of opening by duality. For example, since opening removes the narrow parts of the object and deforms convex corners of the object, we can deduce that closing fills the narrow parts of the background and deforms concave corners of the object. This duality is illustrated by the following figure.

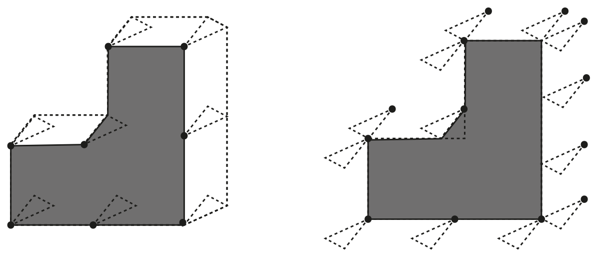

Figure 2.4: An illustration of the duality of opening and closing. Left is an opening of the image, while right is a closing of the background.

As with erosion and dilation, we list some of the useful and important properties concerning the two operations:

- Concerning extensiveness:

  The opening by P is anti-extensive ($\gamma_P(X) \subseteq X$)
  The closing by P is extensive ($X \subseteq \varphi_P(X)$)

- The opening and closing are increasing with respect to the image:

$$X \subseteq Y \Rightarrow \gamma_P(X) \subseteq \gamma_P(Y)$$
$$X \subseteq Y \Rightarrow \varphi_P(X) \subseteq \varphi_P(Y)$$

One of the most notable and useful properties about the opening and closing operators is idempotency:

Proposition 2.1.3. *The opening and closing are idempotent:*

$$\gamma_P \circ \gamma_P = \gamma_P \ and \ \varphi_P \circ \varphi_P = \varphi_P$$

This is particularly significant in the field of image processing where morphology was first applied, because it means that these operators extract all the information available to them after only one application. This means that they are morphological filters, i.e. increasing and idempotent operators.

2.1.3 Summary

These four operators constitute the basis on which mathematical morphology is laid upon. New filters can be obtained by composing closing and opening with the same structuring element, which give four new idempotent filters (opening followed by closing, closing followed by opening, opening followed by closing and then by opening, and finally closing followed by opening and then by closing). Iterative methods on these operators can also be used in order to introduce new filters.

Among those four fundamental operators, erosion and opening are said to be *analytic* and dilation and closing are said to be *generative*. This is due to the former being anti-extensive and the latter being extensive (under the condition $O_E \in P$ for dilation and erosion). As such, erosion and opening are capable of extracting information from a dataset, which makes them prime tools for the analysis of music. On the other hand, dilation and closing enrich a dataset and create more information, which can be useful to reveal the proximity between parts of a dataset.

Through duality, composition and idempotency, the morphological operations are strongly linked between one another, creating a robust tool set. The relationships between the operators is summarized in the figure below:

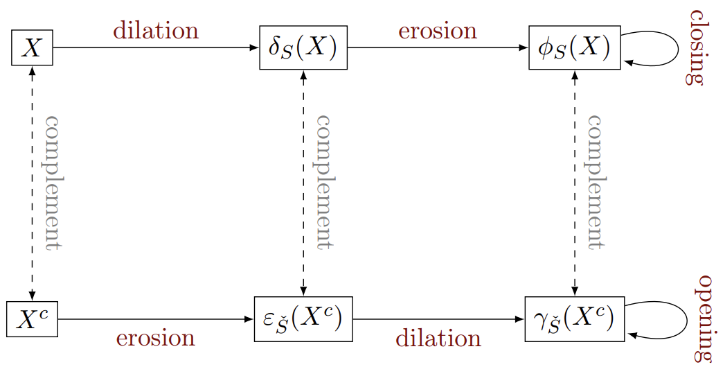

Figure 2.5: The links between the four main operations of mathematical morphology (Source: [13]).

2.2 Point-set Algorithms

The point-set algorithms presented in this section were developed with the intent of finding musically interesting patterns as well as their occurrences. As stated in our introduction, they were developed as an alternative to string-based approaches that were ill-adapted for unvoiced polyphonic music or pieces that were using embellishments.

What constitutes their specificity is that they are based on a multidimensional representation of music. Here the musical data are represented as a set $\mathbb{R}^n$, where each point represents a musical note and each dimension a property of the notes. Generally, we use the space $\mathbb{R}^2$ to represent the pitch and onset of the notes, but other parameters can be considered to enrich the analysis (e.g. the duration, or the voice which the note is part of in the case of voiced music). This multi-dimensional representation of music is close to traditional music representations, such as the score or the MIDI format.

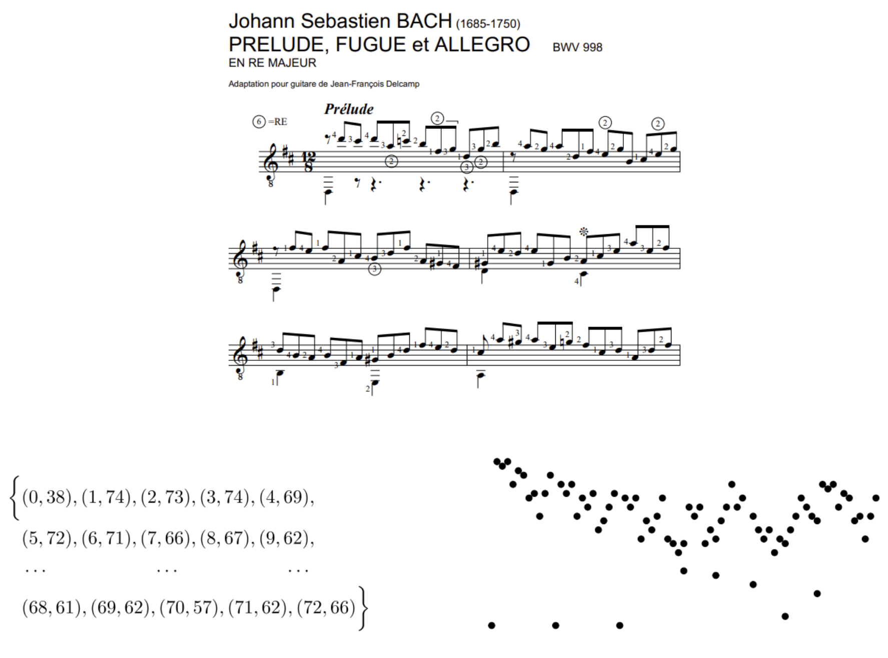

Figure 2.6: Bach's Prelude in E flat major, BWV 998 represented in three different ways: sheet music (top), 2-dimensional point set (bottom-right) and MIDI representation (bottom-left). Source: [14].

In the following, we use X to refer to the dataset we wish to analyze, and P to refer to a pattern.

2.2.1 Maximal Translatable Patterns and the SIA Algorithm

As a first approach, musically interesting patterns in a dataset representing a musical piece can be thought of as *Maximal Translatable Patterns*, otherwise called *MTPs*.

Definition 2.2.1. *Given a vector $v \in E$, the* ***MTP*** *for v in X is defined as follows:*

$$MTP(v, X) = \{x \in X \mid x + v \in X\}$$

Indeed, a pattern can be viewed as a subset of points from the dataset that can be translated within the dataset. By choosing maximal translatable patterns, i.e. the largest pattern that can be translated according to a vector, we ensure that no point that could be part of a musically interesting pattern is left out. Of course, in practice and as we have mentioned in the introduction, other kind of transformations can be used to create repeating patterns, but the restriction to only translated patterns is a good starting point as it is the clearest example of what a pattern can be.

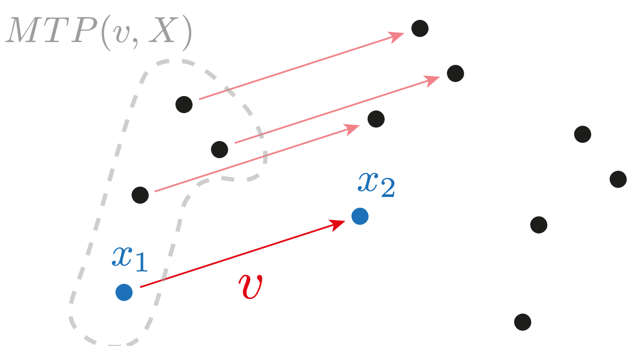

Figure 2.7: Result of $MTP(v, X)$ where X designates the point-set and $v = x_2 - x_1$.

The **Structure Induction Algorithm** [8], also called SIA, is an algorithm designed to compute all the MTPs of a dataset and constitutes the basis for point-set algorithms with the aim of discovering musical patterns.

SIA is designed in such a way that it computes all the translations that can be found within a dataset and the MTPs for these translations. If we define an ordering < on E, and denote a couple of objects A and B with the notation $< A, B >$, then the algorithm returns the following set.

$$\text{SIA: } \{< x_2 - x_1, MTP(x_2 - x_1, X) > \ | \, x_1, x_2 \in X \wedge x_1 < x_2\}$$

This is done in a worst-case running time of $O(k.n^2 log_2 n)$ for a k-dimensional dataset of size n [8]. This algorithm is thus capable of computing translatable patterns within a dataset. Recall however that our objective is not only to find patterns, but also their occurrences. This is where the concept of *Translation Equivalence Class*, or *TEC*, enters the fray.

Definition 2.2.2. *Let $P \in \mathcal{P}(E)$, the* ***TEC*** *for P in X is defined as follows:*

$$TEC(P, X) = \{Q \in \mathcal{P}(E) \,|\, \exists t \in E \; s.t. \; P_t = Q \wedge Q \subseteq X\}$$

Definition 2.2.3. *Let $P, X \in \mathcal{P}(E)$, the* ***set of translators*** *of P in X is defined as follows:*

$$T(P, X) = \{t \in E \,|P_t \subseteq X\}$$

The TEC of a certain pattern can be represented by the pattern and its set of translators, instead of representing each point that is part of an element of the TEC. Doing so is generally less costly when it comes to data-storage.

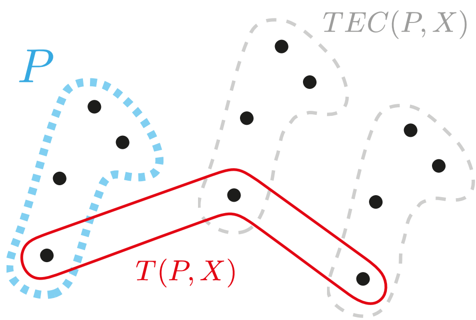

Figure 2.8: The TEC and set of translators of P where the first point of P is the origin of E.

The SIATEC algorithm [8], uses all the patterns discovered with SIA and computes the TECs associated with these patterns. In other words, SIATEC computes the largest translatable patterns P in X and the set of translators of P in X (since a TEC can be represented by a pattern and its set of translators).

SIATEC:

$$\{< MTP(x_2 - x_1, X), T(MTP(x_2 - x_1, X), X) > \mid x_1, x_2 \in X \wedge x_1 < x_2\}$$

The two algorithms we have just described constitute the basis for point-set algorithms for discovering musical patterns. However, they present a few glaring problems:

- most patterns discovered are not musically interesting;
- **the problem of isolated membership:** a musically significant pattern can be contained within a MTP along with other temporally isolated members;

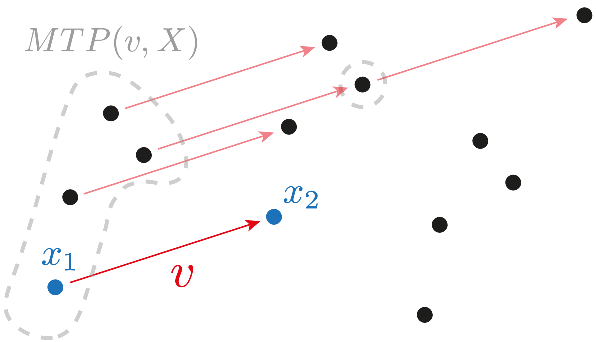

Figure 2.9: Illustration of the problem of isolated membership: one point that can be translated with a musical motif is clearly not part of it.

- the algorithm may not discover musically interesting patterns that occur with variations.

While the last point is the subject of this thesis, the two other problem can generally be dealt with by using heuristics to select musically interesting patterns.

2.2.2 The Use of Heuristics to Select Patterns

In order to select musically relevant patterns we need to define a set of heuristics that allows us to discriminate the good contenders from the bad ones. In this section we look at some of these heuristics and present some of the algorithms derived from SIA that use them. We first introduce compactness.

To define compactness, we first need to define what the *region* of a set of points is. This can be defined in a variety of ways, e.g. with a time segment, a bounding box or a convex hull. One can thus define a region function R that associates to any set of points a subset of E.

Definition 2.2.4. *Given a region function R, the **compactness** C(P,X) of a pattern P in X is defined as the ratio of the number of points of P to the number of points of X in $R(P)$*

Musically significant patterns generally have more compact occurrences: notes in a group that constitutes a coherent whole are generally close to each other. Figure 2.10 shows different how differently defined regions of a pattern lead to different compactness values.

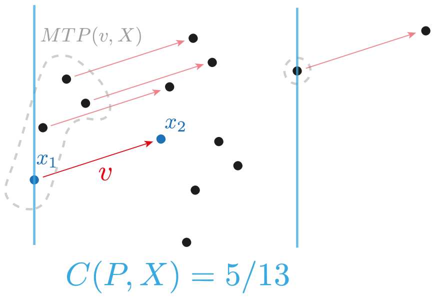

(a) Time segment

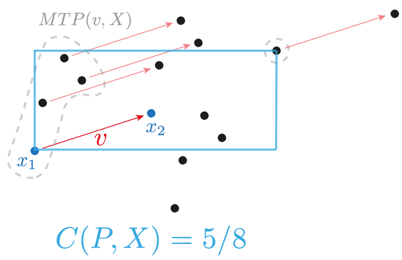

(b) Bounding Box

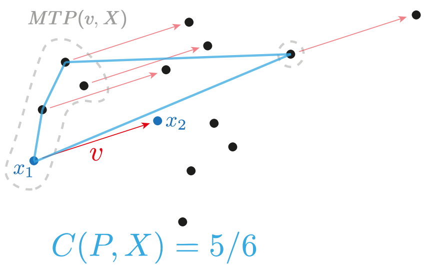

(c) Convex Hull

Figure 2.10: Compactness of a pattern in differently defined regions.

This notion of compactness can be used to deal with the problem of isolated membership. It is notably used in the SIACT algorithm [15]. The aim of SIACT is, for P a pattern found by SIA, to avoid having a pattern $P' \subset P$ that is musically more important than P. In order to do that, SIACT applies a "compactness trawler" to the MTPs discovered by SIA. More precisely, it computes the compactness of MTP sub-patterns to determine whether they should be considered as a pattern. The algorithm is given two threshold values: a compactness threshold and a point threshold. Sub-patterns not satisfying these two thresholds are then deleted.

Another useful concept for finding musically interesting patterns is that of the *compression ratio*. To understand what it is we first need to define the *covered set* of a TEC in a dataset.

Definition 2.2.5. *Let $P, X \in \mathcal{P}(E)$ and $TEC = TEC(P, X)$, the* ***covered set*** *$COV(TEC, X)$ of TEC in X is defined as follows:*

$$COV(TEC, X) = \bigcup_{P \in TEC} P$$

From this concept, we can define the compression ratio of a TEC.

Definition 2.2.6. *Let $P, X \in \mathcal{P}(E)$ and $TEC = TEC(P, X)$, the* ***compression ratio*** *$CR(TEC, X)$ of TEC in X is defined as follows:*

$$CR(TEC, X) = \frac{|COV(TEC,X)|}{|P|+|TEC(P,X)|-1}$$

The compression ratio represents the data storage saved when representing the points of the covering of a pattern by the pattern and its translator set. A higher compression ratio emphasizes patterns that do not overlap and that repeat a lot, which is generally the case for musically relevant patterns. Therefore, a good heuristic for selecting musical patterns is to prioritize patterns with a high compression ratio.

The COSIATEC algorithm, which stands for "COmpression with SIATEC", uses this heuristic to great effect. Its aim is to cover the entire dataset with a set of relevant non-overlapping TECs [16].

First the SIATEC algorithm is applied to obtain the set of all TECs of X. The algorithm then selects the "best" TEC (highest compression ratio). Next all the points covered by the "best" TEC are removed from the dataset, and from then on we repeat the algorithm loop. The algorithm stops when X is empty. COSIATEC produces a family of TECs that covers the entire dataset without overlap.

The main problem with COSIATEC is that it is costly in time: it calls repeatedly SIATEC, which has already a running time cost of $O(n^3)$. SIATECCompress is a less costly alternative. Its aim is also to produce a set of TECs covering the X dataset, but this time the TECs may overlap [16].

SIATECCompress runs SIATEC just once, and then ranks the TECs in order from the best to the worst compression ratio, selecting TECs to the result set by taking a trip down the list and adding the TECs that introduce enough points to be deemed interesting.

2.2.3 MTECs: Another Class of Repeating Patterns

In a 2013 article, T. Collins and D. Meredith present the concept of maximal TEC, otherwise called MTEC [17].

Definition 2.2.7. *Let $P, X \in \mathcal{P}(E)$ and $|TEC(P, X)| = m$. P is a **MTEC** if:*

$$\forall P' \in \mathcal{P}(E), P \subsetneq P' \Rightarrow |TEC(P', X)| < m$$

In other words, if an MTEC is a pattern such that when it is enriched with new points, then the number of new pattern occurrences necessarily decreases. To understand the usefulness of this definition, we may want to look at a characterization of MTECs.

Proposition 2.2.1. *Let $P \subset X$:*

$$P \text{ is } \textbf{MTEC} \Leftrightarrow P = \bigcap_{v \in T(P,X)} MTP(v, X) \text{ [13]}$$

From this characterization, one can show that MTPs are themselves MTECs. Therefore, MTECs can be viewed as a generalization of MTPs, where a pattern is defined by the fact that it can be translated by one or several vectors. Notably, the concept of MTECs can help us to avoid the problem of isolated membership, as illustrated in Figure 2.11.

One of the main difficulties in using MTECs to find musically relevant patterns is that it is unpractical to compute all the MTECs of a dataset. This class of patterns is far bigger than the the class of MTPs, and computing them all would require too big of a complexity. And even if we could reasonably compute all MTECs of a dataset, most of the found patterns would be musically uninteresting just like MTPs. Thus in order to understand how MTECs can be used efficiently to discover patterns, we need another angle of attack. In Section 2.3, we explore how morphology can be useful in order to find MTECs.

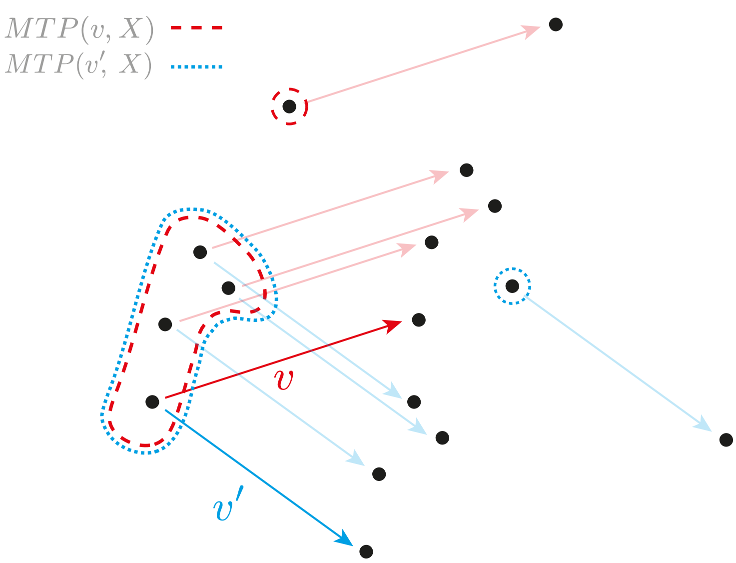

Figure 2.11: Illustration of the characterization of MTECs: the four points that are both in the blue and red envelopes form an MTEC that is the intersection of two MTPs.

2.2.4 Transformable Patterns

As stated in our introduction, translation is not the only way that a pattern can be transformed. In a recent article [18], D. Meredith broadens the notion of MTP to encompass other types of transformations. We first need a few definitions.

Definition 2.2.8. *A* ***transformation class*** F *on* E *is a family of bijections from* E *to* E.

Definition 2.2.9. *Let* $P \subseteq X$. P *is said to be* ***transformable*** *in* X *with respect to a transformation class* F *if there exists a transformation* $f \in F$ *such that* $f(P) \subseteq X$.

We can then generalize the notion of TEC to other transformations, as well as the covered set that is derived from it:

Definition 2.2.10. *Let* $P \subseteq X$. *The* ***occurrence set*** *of* P *in* X *with respect to* F *is defined as:*

$$OS(P, X, F) = \{P\} \cup \{f(P) \mid f(P) \subseteq X \text{ and } f \in F\}$$

Definition 2.2.11. *Let $P \subseteq X$ and $S = OS(P, X, F)$. The* ***covered set*** *of S is defined as:*

$$COV(S) = \bigcup_{P \in S} P$$

Finally, we can define the Maximal Transformable Pattern (MTFP), which generalizes the concept of MTP:

Definition 2.2.12. *Let $f \in F$ a transformation. The* ***maximal transformable pattern*** *for f in X is:*

$$MTPF(X, f) = X \cap f^{-1}(X)$$

This corresponds to the set of points in X that are mapped to other points of X by f. Note that when f is a translation by a vector v, we find that this definition coincides with that of the MTP for v in X. The following property is also worth noting:

Proposition 2.2.2. *Let $f \in F$ a transformation. We then have:*

$$MTFP(X, f^{-1}) = f(MTFP(X, f)) \text{ [14]}$$

Using these concepts, D. Meredith proposes an algorithm capable of computing the maximal transformable patterns of size greater than a certain number. It generalizes the strategy adopted in SIA to any user-specified transformation class over the data space E. Thus, when the transformation class is a class of translations, the algorithm reduces to being SIA.

2.3 Morphology as an Underlying Structure of Point-Set Algorithms

2.3.1 Point-set Theory can be Reformulated using Morphological Operations

P. Lascabettes showcases the link between mathematical morphology and the geometric approach at the core of point-set algorithms [13], allowing for a better understanding of what they entail. More precisely, he showed that most of the concepts used for point-set algorithms can be redefined in morphological terms, from the maximal translatable patterns to the translational equivalence pattern including the covered set. His main results are as follows.

Proposition 2.3.1. *Let $x_1, x_2 \in X$ and $\{x_1, x_2\}$ the structuring element with origin set on x_1. We have:*

$$MTP(x_2 - x_1, X) = \varepsilon_{\{x_1,x_2\}}(X)$$

Proposition 2.3.2. *Let $P, X \in \mathcal{P}(E)$, we have:*

$$TEC(P, X) = \{\delta_t(P) \in \mathcal{P}(E) \mid t \in \varepsilon_P(X)\}$$

Proposition 2.3.3. *Let $P, X \in \mathcal{P}(E)$, we have*

$$T(P, X) = \varepsilon_P(X)$$

Proposition 2.3.4. *Let $P, X \in \mathcal{P}(E)$ and $T = TEC(P, X)$, we have:*

$$COV(T, X) = \gamma_P(X)$$

However, the most important results concerns MTECs, and the way they can be characterized through morphological erosion. We must here mention two important lemmas.

Lemma 2.3.1. *Let $P, X \in \mathcal{P}(E)$, we have:*

$$P \subset \varepsilon_{\varepsilon_P(X)}(X)$$

Lemma 2.3.2. *Let $P, X \in \mathcal{P}(E)$, we have:*

$$\varepsilon_P(X) = \varepsilon_{\varepsilon_{\varepsilon_P(X)}}(X)$$

Those lemmas shed light on a fundamental property of morphological erosion, namely that it involves a looping behavior with respect to the structuring element (see Figure 2.14). If one erodes an image X by a certain set of points P, giving the set $O = \varepsilon_P(X)$, then by eroding X by O, it is ensured that P is contained in $\varepsilon_O(X)$. Moreover, by eroding once again the space by $\varepsilon_O(X)$, one obtains the original erosion O.

These lemmas can then be used to prove the following theorem and its corollary [13].

Theorem 2.3.1. *Let $P, X \in \mathcal{P}(E)$, we have:*

$$P \text{ is MTEC} \Leftrightarrow P = \varepsilon_{\varepsilon_P(X)}(X)$$

Corollary 2.3.1.1. *Let $P, X \in \mathcal{P}(E)$, we have:*

$$P \text{ is MTEC} \Leftrightarrow \exists S, P = \varepsilon_S(X)$$

This significance of this result must be emphasized, as not only it gives us a characterization of MTECs in morphological terms, but according to the corollary we can affirm that the MTECs of a dataset are the patterns discovered by erosion. This result gives us an angle with which we can devise methods in order to find musically relevant MTECs. Mathematical morphology thus gives us a way of discovering musical patterns that are not MTPs (see Fig 2.12).

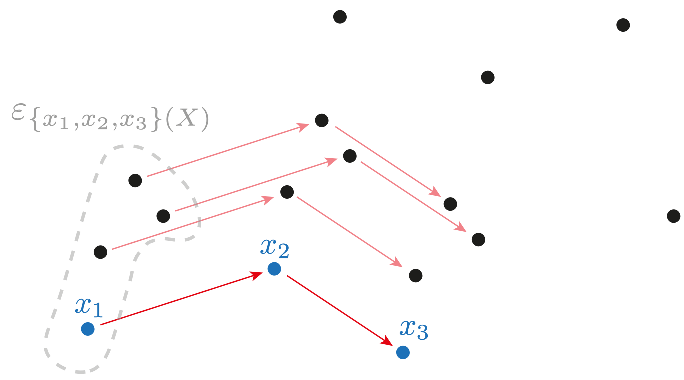

Figure 2.12: An MTEC pattern obtained with the erosion of the dataset X by the set $S = \{x_1, x_2, x_3\}$.

2.3.2 Using Morphology to Discover Musical Patterns

One way of approaching the problem of discovering a musical pattern through mathematical morphology is to try to find it from its *onsets*:

Definition 2.3.1. *Let $P, X \in \mathcal{P}(E)$, the* ***onsets*** *of P in X are defined by:*

$$O = \varepsilon_P(X)$$

The onsets of a pattern can be viewed as the starting points of each of its occurrences (if the origin of the pattern is placed on its first note). With the notations of the above definition, what we call the problem of discovering a pattern from its onsets is understanding when the following equation is true:

$$P = \varepsilon_O(X)$$

The reason why we are interested in this problem is because it is generally easier to discover the onsets than the patterns. This is due to the temporal nature of music, and more specifically to the fact that generally any note (or rest) value is defined as a multiple of certain time span in a music piece, which often induces periodicity when repeating patterns.

Figures 2.13 illustrates a case where we can indeed discover a pattern from its onsets.

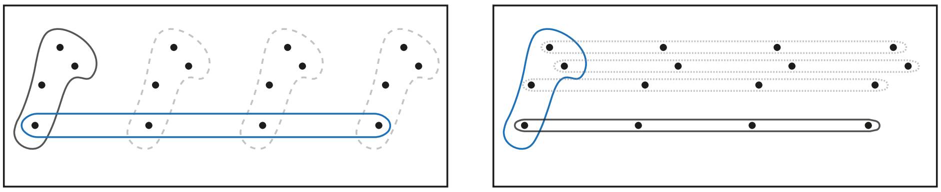

(a) A repeating pattern and its onsets

(b) The pattern is rediscovered when looking at the onsets of the onsets

Figure 2.13: Discovering a pattern from its onsets?

However, in order to understand the depth of the relationship between patterns and onsets, we need to look at the concept of MTEC conjugate pairs.

Definition 2.3.2. *Let $P, O, X \in \mathcal{P}(E)$. The pair (P, O) is an MTEC conjugate pair if it satisfies:*

$$P = \varepsilon_O(X) \text{ and } O = \varepsilon_P(X)$$

MTEC conjugate pairs correspond to patterns and their onsets that can be deduced from each other. Once we have found such a pair, we have found the complete information on a pattern in the data: using the erosion again just leads us in a loop.

Given a musical pattern P, it cannot always be obtained from its onsets O with an erosion, since several different patterns can have the same onsets and the erosion of these onsets can only be equal to one of these onsets at most. However, the following theorem ensures that there exists an MTEC couple (P', O) such that the onsets of P' are the same as those of P:

Theorem 2.3.2. *Let $P, O, X \in \mathcal{P}(E)$ such that $O = \varepsilon_P(X)$ and $O \neq \emptyset$. By defining $P' = \varepsilon_O(X)$, we have:*

$$O = \varepsilon_{P'}(X)$$

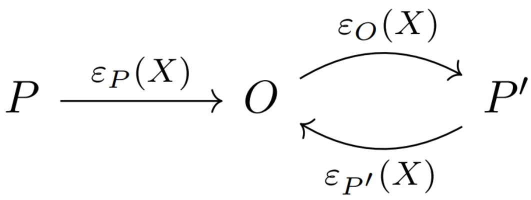

Figure 2.14: The looping behavior of erosion.

What's more, Lemma 2.3.1 ensures that $P \subseteq P'$. Thus, given a musical pattern, we can always complete it with additional notes to obtain an MTEC conjugate pair of a pattern and its onset.

Note that formally speaking, patterns and onsets play a similar role, which means that we can also use the symmetrical approach: starting from a set of onsets, we can complete them to obtain an MTEC conjugate pair of a pattern and its onset.

2.3.3 Discovering Musical Patterns with Variations with Variational Operations

In his master's thesis [14], E. Tamagna proposes an approach to discover musical patterns with variations, which is precisely the topic we are interested in. His approach is based on the definition of an operation akin to morphological erosion, but that allows for the discovery of patterns with a certain degree of similarity. To do so, we first need to define an application that can associate to a pattern other patterns that are similar in a certain sense.

Definition 2.3.3. *A* ***variational class application*** *is an application* $Cl : \mathcal{P}(E) \to \mathcal{P}(\mathcal{P}(E))$ *that associates to a pattern* P *a family* $Cl(P) \subset \mathcal{P}(E)$ *such that* $P \in Cl(P)$*. We say that* $Cl(P)$ *is a variation class for* P*.*

This allows us to define what we can call a *variational erosion*, which plays a role similar to binary erosion but accounts for variations of the pattern as defined by the variational class application.

Definition 2.3.4. *We call the* ***variational erosion*** *by* P *in* X *with respect to* Cl

$$\tilde{\varepsilon}_{Cl(P)}(X) = \{x \in E \mid \exists Q \in Cl(P), Q + x \subset X\}$$

After having defined a form of "variational morphology", one can generalize some of the concepts at the core of point-set algorithms.

Definition 2.3.5. *Let $P \in \mathcal{P}(E)$. The* ***Variational Equivalence Class*** *(or* ***VEC****) of P in X with respect to Cl is defined as:*

$$VEC(P, X, Cl) = \{R \subset X \mid \exists Q \in Cl(P), \exists t \in E, Q_t = R\}.$$

Definition 2.3.6. *Let $P \in \mathcal{P}(E)$. We say that P is a* ***Maximal Variational Equivalence Class*** *(or* ***MVEC*** *in X with respect to Cl if:*

$$\forall P' \in \mathcal{P}(E), P \subsetneq P' \implies |\tilde{\varepsilon}_{Cl(P')}| < |\tilde{\varepsilon}_{Cl(P)}|$$

E. Tamagna then proceeds to apply this theory to a specific variational class application, namely the semitone class application. This class application allows us to represent small variations in a pattern, more precisely variations where only one note differs of one semitone compared to the original pattern. The goal was then be to find a generalization to the characterization of MTECs from Theorem 2.3.1, in order to find patterns more easily:

$$P \text{ is MTEC} \Leftrightarrow P = \varepsilon_{\varepsilon_P(X)}(X)$$

However, such a characterization has yet to be found, which limits the practical use of MVECs in an algorithmic context. In the following, we present an alternative approach to the discovery of musical patterns with variations.

Chapter 3

Discovery of Musical Patterns with Variations

In this chapter, we present the mathematical concepts developed in order to find the variations of patterns and the results ensuring the validity of our approach. More specifically, this chapter is dedicated to morphological methods and operations that can help us finding patterns with their variations.

Figure 3.1 presents a typical case. Given a certain pattern P, we can see that morphological erosion of a data set by P gives us the occurrences of this pattern. However, an occurrence of the pattern that is slightly altered is not picked up by this operation. What we aim to do then is defining operations that allow us to find such variations.

We mainly present two approaches to finding variations, based on two operations that we call *Opening-less Variational Erosion* and *Pivotal Variational Erosion*. Both approaches present some benefits and limitations, be it from a mathematical or musical perspective. We also present another operation called *Intersectional Variational Erosion*, that also presents interesting properties.

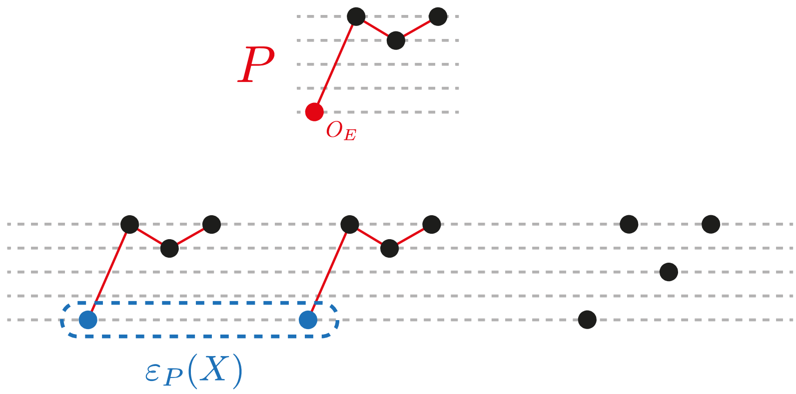

Figure 3.1: The limitations of morphological erosion. When applying the erosion of a dataset X by a pattern P, slightly altered occurrences of P are not detected.

3.1 Opening-less Variational Erosion

We define the *Opening-less Variational Erosion* in the following:

Definition 3.1.1. *Let* $X, P, A \subset \mathcal{P}(E)$*. The* ***Opening-less Variational Erosion****, or* ***OVE****, of* X *by* P *with respect to* A *is defined by:*

$$\tilde{\varepsilon}_{P,A}(X) = \varepsilon_P(X) \cup \varepsilon_P(\delta_A(X \setminus \gamma_P(X))$$

In practice, and similarly to previous approaches, X represents the dataset we wish to study and P represents a pattern whose occurrences we would like to find. We however give ourselves a *dilating element* A, in order to construct new data points that would allow us to find occurrences of patterns similar to P.

Generally, we impose $O_E \in A$, in order to guarantee $X \setminus \gamma_P(X) \subseteq \delta_A(X \setminus \gamma_P(X))$. Indeed, when dilating the set, we want to keep the original points.

Figure 3.2 shows how the OVE works on the example from Figure 3.1. Given a data set X and a pattern P, we first find the occurrences of P in X through the regular morphological erosion $\varepsilon_P(X)$. We then remove those occurrences, which corresponds to the opening $\gamma_P(X)$, and dilate the remaining points with the chosen dilating element A. Finally, we find occurrences of P in this dilated set with $\varepsilon_P(\delta_A(X \setminus \gamma_P(X))$. The result of the operation is then the union of $\varepsilon_P(X)$ and $\varepsilon_P(\delta_A(X \setminus \gamma_P(X))$.

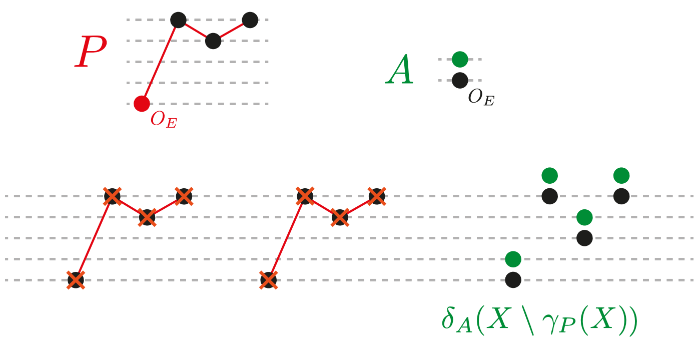

(a) After finding occurrences of P in X, we remove $\gamma_P(X)$ from X and dilate the remaining set of points with A.

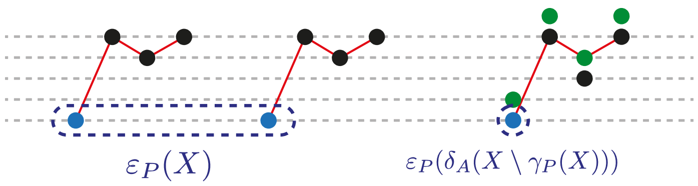

(b) The OVE is the union of $\varepsilon_P(X)$ and $\varepsilon_P(\delta_A(X \setminus \gamma_P(X))$.

Figure 3.2: The OVE applied on a dataset X in order to find a pattern P and its variations.

The choice of A determines the kind of variations we can find. For instance, in the case $E = \mathbb{R}^2$, where the first dimension represents the onsets of notes and the second dimension their pitch, choosing the dilating element $\{(0,0);(0,1)\}$ finds occurrences of a pattern with alterations of pitch, while $\{(0,0);(1,0)\}$ finds occurrences of a pattern with alterations in rhythm. This is illustrated by Figure 3.3.

Moreover, one should note that notes that are not part of the original set X can be found using the OVE (illustrated in Figure 3.3.c). This is not an issue, as such a note, that appeared as a consequence of dilation, can be associated to the original note from which it originated. In the case where the dilating element contains more than two points that are not O_E, it is possible that a constructed note can be associated to several original points. In that case, an order of priority can be defined to decide which point it should be associated with in priority.

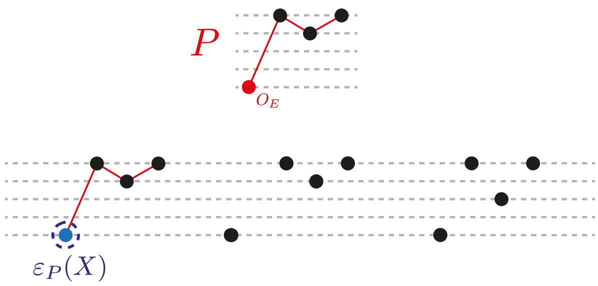

(a) The dataset X contains three occurrences of P: one exact (left), one with variations in rhythm (middle) and one in with variations in pitch (right).

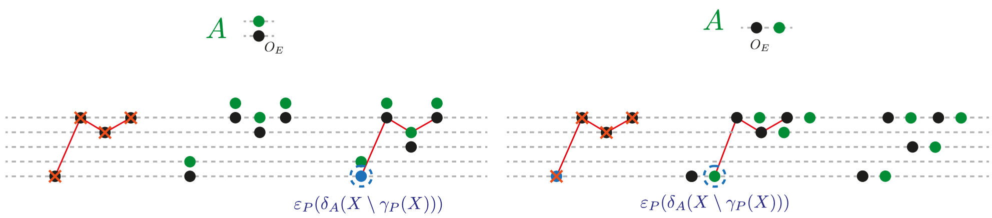

(b) Using $A = \{(0,0);(0,1)\}$, we can find variations of P in pitch.

(c) Using $A = \{(0,0);(1,0)\}$, we can find variations of P in pitch.

Figure 3.3: The impact of the dilating element on the variations of a pattern that are discovered with the OVE.

Recall that in order to use the operation in an algorithmic context, we would like the operation to have a behavior akin to that in Figure 2.14, such that we can find a pattern with its occurrences in a dataset. More precisely, for X a dataset and P a pattern in E, if we define $P' = \tilde{\varepsilon}_{\tilde{\varepsilon}_{P,A}(X),A}(X)$, we would like to have:

$$\tilde{\varepsilon}_{P,A}(X) = \tilde{\varepsilon}_{P',A}(X)$$

This is however generally not the case. Moreover, one can show that that neither $\tilde{\varepsilon}_{P,A}(X) \subseteq \tilde{\varepsilon}_{P',A}(X)$ (see Figure 3.4) nor $\tilde{\varepsilon}_{P,A}(X) \supseteq \tilde{\varepsilon}_{P',A}(X)$ (see Figure 3.5) is generally true.

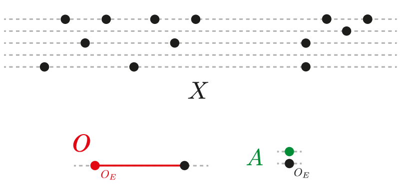

(a) We consider the dataset X, a set of onsets O and a dilating element A.

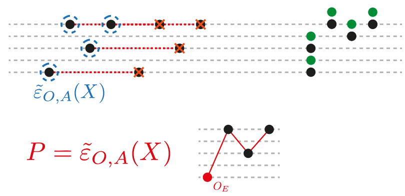

(b) We compute $P = \tilde{\varepsilon}_{O,A}(X)$.

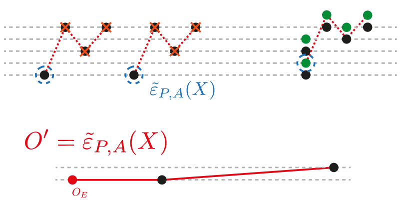

(c) We compute $O' = \tilde{\varepsilon}_{P,A}(X)$.

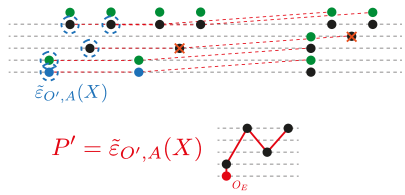

(d) We compute $P' = \tilde{\varepsilon}_{O',A}(X)$. In this case, we have $P \subsetneq P'$.

Figure 3.4: For X a dataset, O a set of onsets and A a dilating element, we have a case where $\tilde{\varepsilon}_{O,A}(X) \subsetneq \tilde{\varepsilon}_{O',A}(X)$ for $O' = \tilde{\varepsilon}_{\tilde{\varepsilon}_{P,A}(X),A}(X)$.

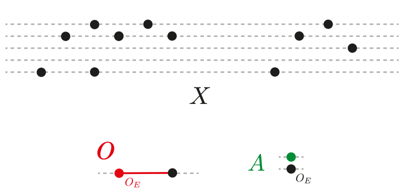

(a) We consider the dataset X, a set of onsets O and a dilating element A.

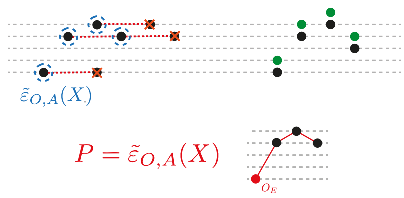

(b) We compute $P = \tilde{\varepsilon}_{O,A}(X)$.

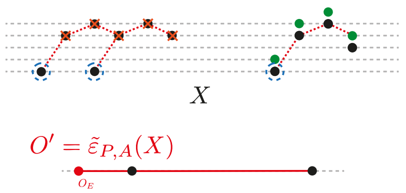

(c) We compute $O' = \tilde{\varepsilon}_{P,A}(X)$.

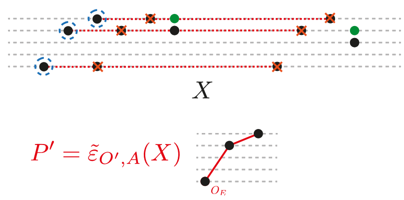

(d) We compute $P' = \tilde{\varepsilon}_{O',A}(X)$. In this case, we have $P \supsetneq P'$.

Figure 3.5: For X a dataset, O a set of onsets and A a dilating element, we have a case where $\tilde{\varepsilon}_{O,A}(X) \supsetneq \tilde{\varepsilon}_{O',A}(X)$ for $O' = \tilde{\varepsilon}_{\tilde{\varepsilon}_{P,A}(X),A}(X)$.

In order to avoid this problem, we define the *OVE with respect to an MTEC conjugate pair*.

Definition 3.1.2. *Let $X, A, Q \in \mathcal{P}(E)$ and (P, O) be an MTEC conjugate pair in X. The* ***OVE with respect to the MTEC conjugate pair*** *(P, O) is defined as a pair of operations:*

$$\tilde{\varepsilon}^{P}_{Q,A}(X) = \varepsilon_P(X) \cup \varepsilon_Q(\delta_A(X \setminus \gamma_P(X)))$$
$$\tilde{\varepsilon}^{O}_{Q,A}(X) = \varepsilon_O(X) \cup \varepsilon_Q(\delta_A(X \setminus \gamma_O(X)))$$

One can note that, due to the properties of MTEC conjugate pairs, we have:

$$\varepsilon_P(X) = O; \varepsilon_O(X) = P; \gamma_O(X) = \gamma_P(X).$$

Therefore, the expression of the OVE with respect to an MTEC couple can be simplified.

Proposition 3.1.1. *Let $X, A \in \mathcal{P}(E)$ and (P, O) be an MTEC conjugate pair in X. By denoting $\gamma = \gamma_O(X) = \gamma_P(X)$, we have for all $Q \in \mathcal{P}(E)$:*

$$\tilde{\varepsilon}^{P}_{Q,A}(X) = O \cup \varepsilon_Q(\delta_A(X \setminus \gamma)$$
$$\tilde{\varepsilon}^{O}_{Q,A}(X) = P \cup \varepsilon_Q(\delta_A(X \setminus \gamma)$$

The idea behind the OVE with respect to an MTEC conjugate pair is to look for variations of patterns that have already been discovered for being part of an MTEC conjugate pair, which makes them candidates for potentially musically interesting patterns. Once we have found a pattern that is part of an MTEC conjugate pair, this set of operations allows us to look for either occurrences of the pattern with alterations, or new points that can be considered part of the pattern.

Theorem 3.1.1 ensures that we can find a loop with these operations.

Theorem 3.1.1. *Let X,A $\in \mathcal{P}(E)$. Let (P, O) be an MTEC couple in X. If we define:*

$$O' = \tilde{\varepsilon}^{P}_{P,A}(X);$$
$$P' = \tilde{\varepsilon}^{O}_{O',A}(X);$$
$$O'' = \tilde{\varepsilon}^{P}_{P',A}(X).$$

We then have: $O'' = O'$.

Proof. To simplify notations, we can write $X' = \delta_A(X \setminus \gamma_P(X)) = \delta_A(X \setminus \gamma_O(X))$.

We set:

$$\begin{aligned} O' = \tilde{\varepsilon}^{P}_{P,A}(X) = O \cup \varepsilon_P(\delta_A(X \setminus \gamma_P(X))) = O \cup \varepsilon_P(X'); \\ P' = \tilde{\varepsilon}^{O}_{O',A}(X) = P \cup \varepsilon_{O'}(\delta_A(X \setminus \gamma_O(X))) = P \cup \varepsilon_{O'}(X'); \\ O'' = \tilde{\varepsilon}^{P}_{P',A}(X) = O \cup \varepsilon_{P'}(\delta_A(X \setminus \gamma_P(X))) = O \cup \varepsilon_{P'}(X'). \end{aligned}$$

We can show $\varepsilon_P(X') = \varepsilon_{P'}(X')$ by double inclusion.

Firstly:

$$P \subset P' \implies \boxed{\varepsilon_{P'}(X') \subset \varepsilon_P(X')}.$$

Let then $o \in \varepsilon_P(X')$.
Then, by definition of morphological erosion:

$$o \in \varepsilon_P(X') \implies P_o \subset X' \implies \forall p \in P, o + p \in X'.$$

But also, $\varepsilon_P(X') \subset O' \implies o \in O'$.
Thus:

$$\forall p \in \varepsilon_{O'}(X'), O'_p \subset X' \implies o + p \in X'.$$

Since $P' = P \cup \varepsilon_{O'}(X')$, this gives us:

$$\forall p \in P', o + p \in X'.$$

Therefore: $o \in \varepsilon_{P'}(X')$. Thus, we have $\boxed{\varepsilon_P(X') \subset \varepsilon_{P'}(X')}$

Finally:

$$\begin{aligned} \varepsilon_P(X') = \varepsilon_{P'}(X') &\implies O \cup \varepsilon_P(X') = O \cup \varepsilon_{P'}(X') \\ &\implies \boxed{O' = O''} \end{aligned}$$

□

This looping structure gives us a procedure in order to find occurrences with variations of a pattern. Start with an MTEC conjugate pair (P, O) and a dilating element A. Then compute $\tilde{\varepsilon}^{P}_{P,A}(X)$: this gives a set of onsets O' that correspond to occurrences of P and its variations (as defined through A). Finally compute $\tilde{\varepsilon}^{O}_{O',A}(X)$: this gives a set P' that completes the pattern P with new points that can also be considered part of it. Theorem 3.1.1 ensures that this process has brought out all the information available.

Note however that, since both members of an MTEC conjugate pair have formally speaking the same role, one can start this process from O instead of P. Therefore, the procedure can be used in two different ways: to find new occurrences of a pattern (when one starts the procedure with P), or to complete a pattern (when one starts the procedure with O). Figure 3.6 summarizes the looping structure of the procedure.

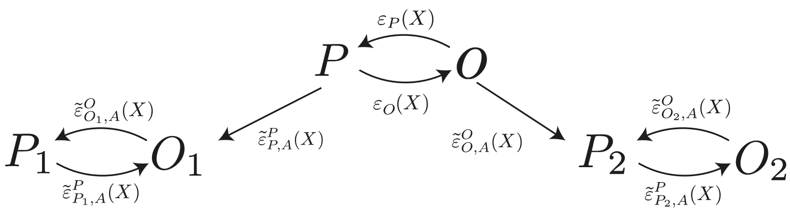

Figure 3.6: The looping behavior of OVE with respect to an MTEC conjugate pair (P, O).

One can wonder whether the patterns P_1 and P_2 and the onsets O_1 and O_2 from Figure 3.6 are related in any way. The following theorem ensures that they are through inclusion.

Theorem 3.1.2. *Let X,A $\in \mathcal{P}(E)$ and let (P, O) be an MTEC conjugate pair in X. If we define (as in Figure 3.6):*

$$\begin{aligned} O_1 &= \tilde{\varepsilon}^P_{P,A}(X); \\ P_1 &= \tilde{\varepsilon}^O_{O_1,A}(X); \\ P_2 &= \tilde{\varepsilon}^O_{O,A}(X); \\ O_2 &= \tilde{\varepsilon}^P_{P_2,A}(X). \end{aligned}$$

Then we have: $P_1 \subset P_2$ and $O_2 \subset O_1$.

Proof. Once again, we can write $X' = X \setminus \gamma_P(X) = X \setminus \gamma_O(X)$, which gives us:

$$O_1 = O \cup \varepsilon_P(X')$$

$$P_1 = P \cup \varepsilon_{O_1}(X')$$

$$P_2 = P \cup \varepsilon_O(X')$$

$$O_2 = O \cup \varepsilon_{P_2}(X')$$

We then have:

$$\begin{aligned} O \subset O_1 &\implies \varepsilon_{O_1}(X') \subset \varepsilon_O(X') \\ &\implies P \cup \varepsilon_{O_1}(X') \subset P \cup \varepsilon_O(X') \\ &\implies P_1 \subset P_2 \end{aligned}$$

By symmetry, we also have $O_2 \subset O_1$. □

Keeping notations from Figure 3.6, it is possible to have $P_1 = P_2$ and $O_1 = O_2$. In that case, the problem of finding new occurrences and the problem of finding new points part of the pattern are equivalent. It is however generally not the case, as illustrated by Figure 3.7.

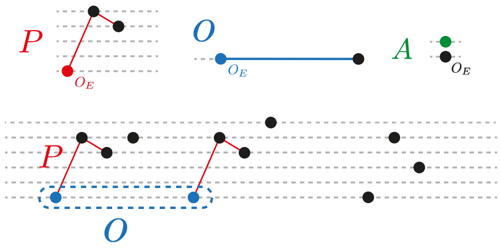

(a) (P, O) is an MTEC conjugate pair in the dataset X.

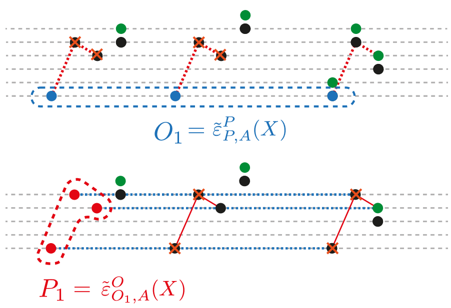

(b) We compute $O_1 = \tilde{\varepsilon}^P_{P,A}(X)$ and $P_1 = \tilde{\varepsilon}^O_{O_1,A}(X)$.

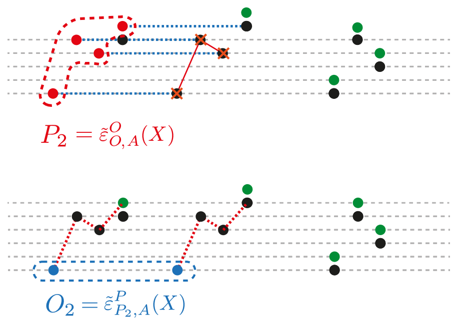

(c) We compute $P_2 = \tilde{\varepsilon}^O_{O,A}(X)$ and $P_1 = \tilde{\varepsilon}^P_{P_2,A}(X)$.

Figure 3.7: From and MTEC conjugate pair (P, O) in a dataset X, we compute O_1,P_1,P_2 and O_2 like in Figure 3.6. We can see that $P_1 \subsetneq P_2$ and $O_1 \supsetneq O_2$. O_1 corresponds to new occurrences of the pattern and P_2 to new points of the pattern found with the OVE with respect to (P, O)

The main drawback with this approach is that it forces us to work from MTEC conjugate pairs in order to find patterns with variations in a dataset. This is a limitation, as it prevents us from discovering new patterns when compared to the regular approach, and only allows us to get more information on an already discovered pattern. We can find new occurrences of the pattern, or complete it with new points, but not discover a whole new pattern.

However, one should recall that more is not necessarily better when it comes to finding musically interesting patterns in a music piece. Sometimes, finding less patterns can add clarity to the analysis of the piece. Thus, getting more information on already found MTEC conjugate pairs is a good way to ensure that the new information is meaningful. Notably, it sometimes allows us to determine that two different MTEC conjugate pairs actually represent the same pattern when variations are accounted for.

3.2 Pivotal Variational Erosion

We define the *Pivotal Variational Erosion* in the following manner:

Definition 3.2.1. *Let* $X, P, A \subset \mathcal{P}(E)$ *and* $x \in E$*. The* ***Pivotal Variational Erosion****, or* ***PVE****, of* X *by* P *with respect to* A *is defined by:*

$$\hat{\varepsilon}^{x}_{P,A}(X)_{=}\{t \in E \mid P_t \subseteq \delta_A(X) \wedge t + x \in X\}$$

In practice, just like in the case of the OVE, X is the dataset we wish to study, P a pattern whose instances we want to find in X, and A a dilating element that defines the types of variations we want to find. We however have an additional parameter x, called the *pivot*, that allows us to only keep points that are part of the original dataset X after translation by x.

Once again, we generally impose $O_E \in A$, so that $X \subset \delta_A(X)$. This time, we also generally choose A to be symmetric with respect to O_E. Doing so ensures that variations are kept whether they occur in one sense or the other. Finally, we also generally choose $x \in P$, in order to look for instances of P in the dilated dataset with at least one point in common with the original dataset.

Figure 3.8 illustrates how this operations functions on the example from Figure 3.1.

We can show that this operation has a looping behavior with well chosen pivot points. First, the following lemma ensures that the PVE is *decreasing* with respect to the structuring element.

Lemma 3.2.1. *Let* $X \in \mathcal{P}(E)$*,* $A \in \mathcal{P}(E)$ *and* $x \in E$*. Let* $P, P' \in \mathcal{P}(E)$ *such that* $P \subseteq P'$*. Then:*

$$\hat{\varepsilon}^{x}_{P',A}(X) \subseteq \hat{\varepsilon}^{x}_{P,A}(X).$$

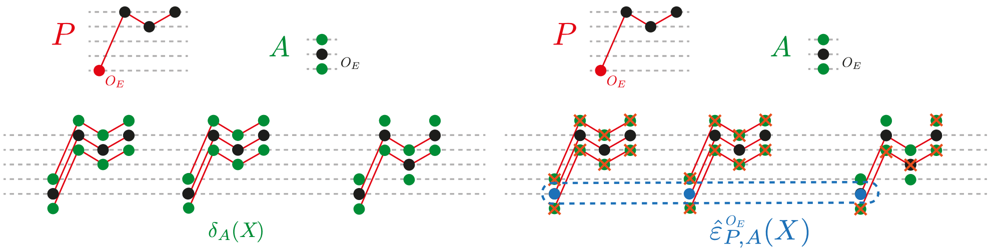

(a) We look for instances of a pattern P in the dilation of the dataset $\delta_A(X)$.

(b) We keep those instances if the first point of P is part of the original dataset X.

Figure 3.8: The PVE applied on a dataset X with pivot O_E, which is the leftmost point of P, in order to find a pattern P and its variations.

Proof. Let $t \in \hat{\varepsilon}^{x}_{P',A}(X)$. Then:

$$(P'_t \subseteq \delta_A(X) \wedge P_t \subseteq P'_t) \implies P_t \subseteq \delta_A(X)$$

Also: $t \in \hat{\varepsilon}^{x}_{P',A}(X) \implies t + x \in X$.
Therefore:

$$(P_t \subseteq \delta_A(X) \wedge t + x \in X) \iff t \in \hat{\varepsilon}^{x}_{P,A}(X)$$

□

This second lemma then gives us an inclusion relationship between a pattern and the PVE of the PVE of the pattern (with well chosen pivots), provided the pattern can be found in the original dataset. This is a generalization of Lemma 2.3.1 for variations.

Lemma 3.2.2. *Let $X, A \in \mathcal{P}(E)$ and $x \in E$. Let $P \in \mathcal{P}(E)$ and $t \in E$ such that $P_t \subset X$. Then:*

$$P \subseteq \hat{\varepsilon}^{t}_{\hat{\varepsilon}^{x}_{P,A}(X),A}(X).$$

Proof. Let us write $O = \hat{\varepsilon}^{x}_{P,A}(X)$ and let $p \in P$.
We need to show:

$$p \in \hat{\varepsilon}^{t}_{O,A}(X)$$

In other words, we want to show:

$$\overbrace{O_p \subseteq \delta_A(X)}^{(i)} \wedge \overbrace{p + t \in X}^{(ii)}$$

Firstly, we have $P_t \subset X$. Therefore: $\boxed{p + t \in X}$ (ii)
Then, by definition of O:

$$\begin{aligned} \forall o \in O, P_o \subseteq \delta_A(X) &\implies \forall o \in O, \forall q \in P, q + o \in \delta_A(X) \\ &\implies \forall q \in P, O_q \subseteq \delta_A(X) \\ &\implies \boxed{O_p \subseteq \delta_A(X)}\ (i) \end{aligned}$$

Thus: $\boxed{p \in \hat{\varepsilon}^t_{O,A}(X)}$ □

Using those two lemmas, we can prove the following theorem that ensures that the PVE has a looping property when choosing appropriate pivot points.

Theorem 3.2.1. *Let* $X, A \in \mathcal{P}(E)$. *Let* $P \in \mathcal{P}(E)$, $x \in P$ *and* $t \in E$ *such that* $P_t \subset X$. *Then:*

$$\hat{\varepsilon}^x_{P,A}(X) = \hat{\varepsilon}^x_{\hat{\varepsilon}^t_{\hat{\varepsilon}^x_{P,A}(X),A}(X)}(X).$$

Proof. We reason by double inclusion.
First, according to Lemma 3.2.2, we have: $P \subseteq \hat{\varepsilon}^t_{\hat{\varepsilon}^x_{P,A}(X),A}(X)$.
Therefore, according to Lemma 3.2.1:

$$\boxed{\hat{\varepsilon}^x_{\hat{\varepsilon}^t_{\hat{\varepsilon}^x_{P,A}(X),A}(X),A}(X) \subseteq \hat{\varepsilon}^x_{P,A}(X)}$$

For the other inclusion let us write $O = \hat{\varepsilon}^x_{P,A}(X)$ and consider $o \in O$.
By definition of O, we have: $o + x \in X$.
Also, if we write $P' = \hat{\varepsilon}^t_{O,A}(X)$:

$$\begin{aligned} \forall p \in P', O_p \subset \delta_A(X) &\implies \forall p \in P', \forall q \in O, q + p \subset \delta_A(X) \\ &\implies \forall q \in O, P'_q \subset \delta_A(X) \\ &\implies P'_o \subset \delta_A(X) \end{aligned}$$

We therefore have:

$$\begin{aligned} \forall o \in O, (P'_o \subset \delta_A(X) \wedge x + o \in X) &\implies \forall o \in O, o \in \hat{\varepsilon}^x_{P',A}(X) \\ &\implies \forall o \in O, o \in \hat{\varepsilon}^x_{\hat{\varepsilon}^t_{O,A}(X),A}(X) \\ &\implies O \subseteq \hat{\varepsilon}^x_{\hat{\varepsilon}^t_{O,A}(X),A}(X) \\ &\implies \boxed{\hat{\varepsilon}^x_{P,A}(X) \subseteq \hat{\varepsilon}^x_{\hat{\varepsilon}^t_{\hat{\varepsilon}^x_{P,A}(X),A}(X),A}(X)} \end{aligned}$$

□

Thus, we have found an operation that loops and discovers patterns and their occurrences with variations defined from an approximation. Considering a dataset X, a set of points P such that $\exists t, P_t \subset X$, and choosing $x \in P$, one can use the PVE $\hat{\varepsilon}^x_{P,A}(X)$ in order to find occurrences of P and its variations as defined by the dilating element A, and then $\hat{\varepsilon}^x_{\hat{\varepsilon}^x_{P,A}(X),A}(X)$ in order to complete P into a "conjugate" of O. One then has found all the information available with this operation. This looping behavior is represented in Figure3.9.

The operation thus presents a clear advantage compared to the previous approach, which is that it does not necessitate to start from an already found MTEC conjugate pair. Thus, it is possible to find patterns unrelated to MTEC conjugate pairs and that appear with variations in a music piece. This approach is thus more powerful than the previous one.

However, as stated before, more is not always better when it comes to finding musically significant patterns. One thus needs to define heuristics to select patterns that might be interesting among the new ones found. Moreover, since this approach dilates the entire dataset and does not remove the opening before dilation in contrast to the previous approach, it is more susceptible to find instances of a pattern that are redundant (e.g. when dilating an instance of the pattern where a unique point is present).

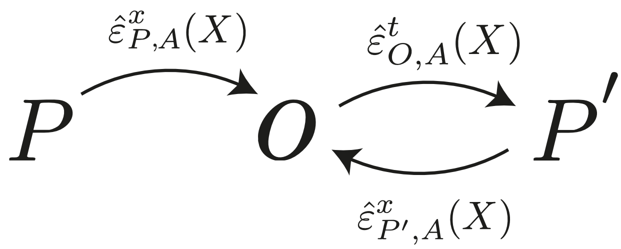

Figure 3.9: The looping behavior of the PVE starting from a pattern P and with $x, t \in E$ such that $x \in P$ and $P_t \subset X$.

3.3 Intersectional Variational Erosion:

We define the *Intersectional Variational Erosion* in the following manner:

Definition 3.3.1. *Let $X, P, A \subset \mathcal{P}(E)$. The* ***Intersectional Variational Erosion****, or* ***IVE****, of X by P with respect to A is defined by:*

$$\bar{\varepsilon}_{P,A}(X) = \{t \in E \mid P_t \subseteq \delta_A(X) \wedge X \cap P_t \neq \varnothing\}$$

In practice, just like in the two previous cases, X is the dataset we wish to study, P a pattern whose instances we want to find in X, and A a dilating element that defines the types of variations we want to find, and we generally impose $O_E \in A$.

Figure 3.10 illustrates how this operations functions.

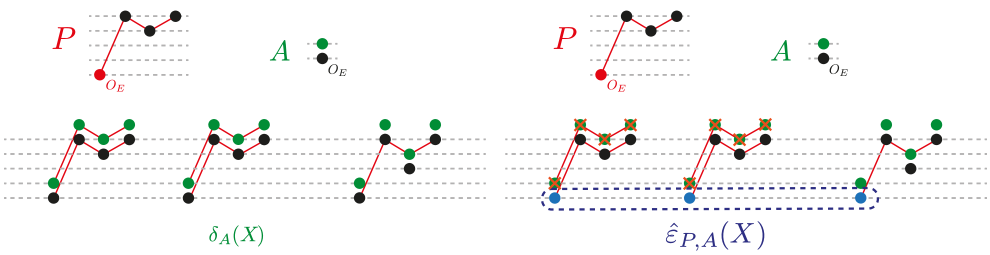

(a) We look for instances of a pattern P in the dilation of the dataset $\delta_A(X)$.

(b) We keep those instances if and only if they have a point in common with the original dataset X.

Figure 3.10: The IVE applied on a dataset X in order to find a pattern P and its variations.

This operation constitutes an example of an operation that appears promising but ultimately does not satisfy the properties we are looking for. It is however close, since it can be guaranteed to loop after a finite number of iterations under certain circumstances. This can be proven with the following lemma.

Lemma 3.3.1. *Let $X \in \mathcal{P}(E)$, $A \in \mathcal{P}(E)$ and $P \in \mathcal{P}(E)$ such that $\exists t, P_t \subset X$ and $O_E \in A$. Then:*

$$P \subseteq \bar{\varepsilon}_{\bar{\varepsilon}_{P,A}(X),A}(X)$$

Proof. Let us write $O = \bar{\varepsilon}_{P,A}(X)$. Let $p \in P$, we need to show:

$$p \in \bar{\varepsilon}_{O,A}(X) \iff O_p \subseteq \delta_A(X) \wedge X \cap O_p \neq \varnothing$$

By definition of O:

$$\begin{aligned}\forall o \in O, P_o \subseteq \delta_A(X) &\implies \forall o \in O, \forall q \in P, q + o \in \delta_A(X)\\ &\implies \forall q \in P, O_q \in \delta_A(X)\\ &\implies \boxed{O_p \in \delta_A(X)}\end{aligned}$$

Also: $O_E \in A \implies X \subset \delta_A(X)$

Therefore: $P_t \subset X \implies \begin{cases} P_t \subset \delta_A(X) \\ X \cap P_t \neq \varnothing \end{cases} \implies t \in O$

Which gives us:

$$\begin{cases} P_t \subset X \\ t \in O \end{cases} \implies \begin{cases} p + t \in X \\ p + t \in O_p \end{cases} \implies \boxed{X \cap O_p \neq \varnothing}$$

Thus: $\boxed{p \in \bar{\varepsilon}_{O,A}(X)}$

□

Using this lemma, one can show that the IVE ends up looping after a finite number of iterations, provided the dataset and dilating element are finite, and the structuring element is part of the dataset.

Theorem 3.3.1. *Let $X \in \mathcal{P}(E)$, $A \in \mathcal{P}(E)$ and $P \in \mathcal{P}(E)$ such that $\exists t, P_t \subset X$ and $O_E \in A$. We define the following series:*

$$(P_n)_{n\in\mathbb{N}} : \begin{cases} P_o = P \\ P_n = \bar{\varepsilon}_{P_{n-1}}(X) \text{ for } n \geq 1 \end{cases}$$

If we suppose that X and A are finite, then: $\exists n \in \mathbb{N}, P_n = P_{n+2}$.

Proof. We reason by induction in order to show: $\forall n \in \mathbb{N}, P_n \subseteq P_{n+2}$.

Firstly, Lemma 3.3.1 ensures that $P_0 \subseteq P_2$

Let us then suppose that $P_{n-1} \subseteq P_{n+1}$ for a certain $n \in \mathbb{N}^*$.
We then have $P_n = \bar{\varepsilon}_{P_{n-1},A}(X)$ and $P_{n+2} = \bar{\varepsilon}_{P_{n+1},A}(X)$.
Let $p \in P_n$, we would like to show:

$$p \in \bar{\varepsilon}_{P_{n+1},A}(X)$$

In other words, we would like to show:

$$(P_{n+1})_o \subseteq \delta_A(X) \wedge X \cap (P_{n+1})_o \neq \varnothing$$

Using the same reasoning as in Lemma 3.3.1, one can show $\boxed{(P_{n+1})_o \subseteq \delta_A(X)}$. Also, by definition of O_{n-1} and by the induction hypothesis, we have:

$$\begin{cases} X \cap (P_{n-1})_o \neq \varnothing \\ P_{n-1} \subseteq P_{n+1} \end{cases} \implies \boxed{X \cap (P_{n+1})_o \neq \varnothing}$$

Thus, we have indeed have $P_n \subseteq P_{n+2}$.

The series $(P_{2n})_{n\in\mathbb{N}}$ is thus increasing.
However, $\forall n \in \mathbb{N}^*$:

$$P_{2n} \subseteq \varepsilon_{P_{2n-1}}(\delta_X(A)) \subseteq \varepsilon_{P_1}(\delta_X(A))$$

Since X and A are supposed finite, then so is $\varepsilon_{P_1}(\delta_X(A))$.
Since $(P_{2n})_{n\in\mathbb{N}}$ is increasing and bounded, $\boxed{\exists n \in \mathbb{N}, P_n = P_{n+2}}$. □

We represent the looping behavior in Figure 3.11. This result makes it possible to use the operation in an algorithmic context, as we can use the operation a finite number of times before finding all the information available. However, doing so can be inefficient, as there is no guarantee that it loops quickly. This is due to the fact that it is not decreasing with respect to the structuring element, as illustrated in Figure 3.12. Not knowing when the process ends is a significant drawback when designing our algorithms, as it potentially increases their complexity.

Moreover, the IVE is very similar to the PVE, the difference being that PVE requires that a specific point of the pattern be part of the original dataset, whereas the IVE only requires that it be the case for any point of the pattern. The looser condition that defines IVE thus makes it even more prone to finding redundancies than the PVE. This can be seen in Figure 3.10.c, where an extra occurrence of the pattern is found due to the proximity between the pattern and an isolated point.

Taking this into account, we favor the use of the PVE rather than the IVE in an algorithmic context.

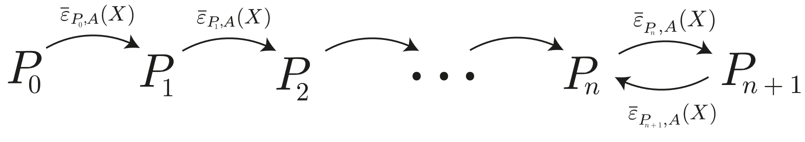

Figure 3.11: The looping behavior of the IVE starting from a pattern $P \subset X$.

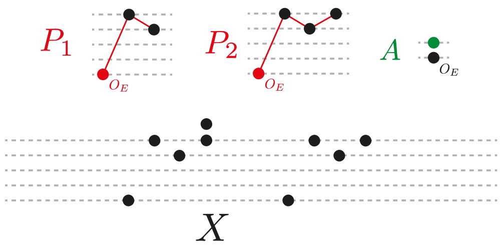

(a) We define a dataset X, a structuring element A and two patterns P_1 and P_2 such that $P_1 \subset P_2$.

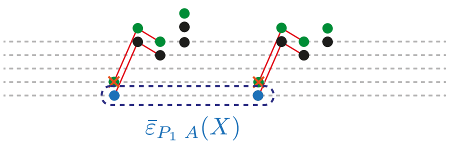

(b) We compute $\bar{\varepsilon}_{P_1,A}(X)$.

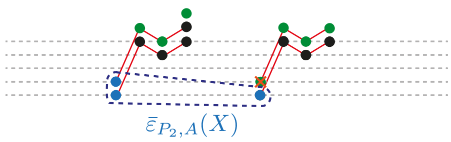

(c) We compute $\bar{\varepsilon}_{P_2,A}(X)$. In this case, we have $\bar{\varepsilon}_{P_1,A}(X) \subsetneq \bar{\varepsilon}_{P_2,A}(X)$

Figure 3.12: The IVE is not decreasing with respect to the structuring element.

Chapter 4

Finding Patterns in Bach's Well-Tempered Clavier

In this section, we apply the mathematical results of the previous section to a concrete case: the Well-Tempered Clavier by J.S Bach. The goal is to find the subject for each of the fugues from the first book of the cycle using the operations defined in the previous chapter. We describe two algorithms: one based on the OVE with respect to an MTEC conjugate pair, and one based on the PVE. Other algorithms for the discovery of musical patterns were used as sources inspiration of [8, 19]. In both cases, we compare the results from the algorithm with those from a paper by M. Giraud, R. Groult and F. Levé [20], in order to determine the strengths and limitations of our approaches.

In the following, one should keep in mind that those algorithms are specifically tailored to the problem of finding the subject of a fugue and its occurrences, and that different strategies should be adopted when it comes to finding unspecified recurring patterns in a piece of music. However, the aim is to demonstrate that using the operations we have defined in Chapter 3 can indeed refine the description of a piece when compared to the use of regular erosion in order to find MTEC conjugate pairs. As algorithms designed to find musically interesting MTEC conjugate pairs improve, they can be combined with our approach. Notably, since MTPs form a subset of MTECs, one could use our approach in the context of the SIA algorithm and its derivatives.

In the following, the notes from the music pieces are represented by points in $E = \mathbb{R}^2$. The first dimension represents the onset of the note in time. We consider that a time value of one quarter note is represented by a value of one. The second dimension represents the pitch in MIDI encoding, e.g. C5 is represented by the value 60. In the following, we use the lexicographic ordering on E:

$$\forall \mathbf{x}, \mathbf{y} \in E, \mathbf{x} < \mathbf{y} \iff (x_1 < y_1) \vee (x_1 = y_1 \wedge x_2 < y_2)$$

4.1 Using the OVE with respect to an MTEC Conjugate Pair

4.1.1 Description

This first algorithm is based on the OVE with respect to an MTEC conjugate pairs. As such, the idea behind this algorithm is to first find musically interesting MTEC conjugate pairs, and then use the operation in order to find new occurrences of the pattern, and possibly new points that could be part of the pattern.

The first step consists in finding MTEC conjugate pairs. Since we are looking for the subject of a fugue, which always appears first during the exposition, we can look for it by considering the initial consecutive notes of the piece. To take into account the cases where there is an anacrusis, we may allow to look for the subject within the consecutive notes after an onset of one or two notes. The process is described by Algorithm 1.

Let X be the dataset we wish to study, and A the dilating element used to study variations. Starting with $n = 1$, we consider the three sets constituted from n consecutive notes, starting from the first, second and third notes. For each set, we consider the erosion of X by each set. If the erosion contains only one element, then it means that the motif appears only once in the piece and thus that it is not a repeating pattern. When this is the case for all three of the considered sets, we can stop the process, as any of the following sets we might consider will contain at least one of those and thus cannot be a repeating pattern. If the erosion contains more than one element, then it means that it is a repeating pattern. We can then erode X once more by the erosion in order to complete the pattern, giving us an MTEC conjugate pair (P, O), with P being a pattern and O its onsets. We then save the pair in a list, if it is not already present in it.

Once this is done, the next step consists in discriminating between patterns that are potentially musically interesting and those that are not. This is where we need to introduce a set of heuristics. We consider two parameters: the number of notes and the compactness of a pattern. The condition on the number of notes is pretty straightforward: we want a pattern to contain sufficiently enough notes to be considered a "subject" of the fugue. The minimum number of notes we decided on is 4. When it pertains to compactness, we first need to define the region of a pattern. To do this, we use the ordering that we have defined on E: the region of a pattern is the set of all points between its first and last points as defined by this ordering. Considering the first occurrence of the pattern, we thus define its compactness as the ratio of its length to the number of points of X that are between its first and last points. The minimum compactness we impose is 0.6. We thus delete any element from the list of MTEC conjugate pairs where the pattern does not satisfy both of these conditions.

Algorithm 1 Finding MTEC conjugate pairs from consecutive notes

```
n ← 1
Pairs ← []
stop ← False
while not stop do
    onsets ← []
    for k ← 1 to 3 do
        S ← the n consecutive notes of X starting from the k-th note
        O ← ε_S(X)
        if |O| > 1 then insert O in onsets
        end if
    end for
    if onsets is empty then
        stop ← True
    else
        for O ∈ onsets do
            P ← ε_O(X)
            if (P, O) not in pairs then
                insert (P, O) in pairs
            end if
        end for
    end if
end while
```

The final step consists in applying the OVA with respect to MTEC conjugate pairs to the MTEC conjugate pairs we have memorized. The process is described in Algorithm 2.

For every MTEC conjugate pairs (P, O) we have memorized, we compute $O' = \hat{\varepsilon}^{P}_{P,A}(X)$ and then $P' = \hat{\varepsilon}^{O}_{O',A}(X)$. We thus have extracted all the information available using the OVE with respect to the (P, O). We then want to memorize the extended pairs (P', O') in a new list. However, one of the interesting aspects of this approach is that it can reveal that MTEC conjugate pairs can be equivalent to one another when extended. The case where this is the clearest is when two MTEC conjugate pairs give way to the same extended pairs, but it needs not necessarily be the case. Recall that notes that are not part of the original set can be found through erosion. However, these notes still refer to notes from the music piece through dilation. Therefore, two patterns can refer to the same motif despite being different. Therefore, in order to avoid redundancies, we compare what one can call the reference of the extended patterns, i.e. the set of points obtained when each point of the extended pattern is associated with the point in the dataset from which it comes through dilation. If their are several such points, one can choose an order of priority. This is referred to as $ref^{A}_{X}(P')$, and is well defined as long as $P' \subset \delta_A(X)$. Thus, for each extended pair computed, we check if the reference of the extended pattern has not already been found before adding it to the list.

Algorithm 2 Using the OVA with respect to MTEC conjugate pairs

```
VarPairs ← []
mem ← []
for all (P, O) ∈ Pairs do
    O' ← ε̂^P_{P,A}(X)
    P' ← ε̂^O_{O',A}(X)
    if ref^A_X(P') is in mem then
        insert (P', O') in VarPairs
        insert P' in mem
    end if
end for
```

Considering an extended couple (P', O') from an MTEC conjugate pair (P, O), using the notion of reference is useful to count the actual number of notes in the pattern and the actual number of occurrences, especially when the structuring element used to dilate the space has more than one non-zero element, which can lead to redundancies. One should however keep in mind that the way the algorithm is laid out implies that $P' \subset \delta_A(X)$ but not necessarily that $O' \subset \delta_A(X)$. When this is the case, the reference of O' can not be properly defined. However, we have $O'_p \subset \delta_A(X)$ for p being the first element of P', meaning that we can compute the reference of O'_p. Thus, $|ref(P')|$ and $|ref(O'_p)|$ give us respectively the actual number of notes in the pattern, and the actual number of occurrences.

In the end, the algorithm gives us a list of extended pairs giving us a pattern and its occurrences in X with variations as defined by the dilating element A. The expected behavior is that the algorithm not only gives us the subject of the fugue and its occurrences with variations, but also recurring partial instances of the subject, which are also interesting to pick up on and can deepen the understanding of the pieces.

The complexity necessary to compute the erosion of a set X by a structuring element P is $O(|P||X| \times ln(|P||X|))$, because it is based on sorting a list of $|P||X|$ elements. To compute an OVA with respect to an MTEC conjugate pair with dilating element A this would give us $O(|A||P||X| \times ln(|A||P||X|))$, as it's complexity is equivalent to that of an erosion on the dilated set. Since we are doing a determined amount of such operations for every set of consecutive notes in X up until a certain point, this gives us a worst case complexity of $O(k \times n^3 ln(k \times n))$, where $k = |A|$ and $n = |X|$.

4.1.2 Results

In this section, we consider a few cases to illustrate some of the benefits and drawbacks of the approach and compare our results with those of M. Giraud, R. Groult and F. Levé [20]. We study the fugues from the first book of the Well-Tempered Clavier using 4 distinct dilating elements, defined below and represented in 1

$$
\begin{aligned}
A_1 &= \{(0,0),(0,1)\};\\
A_2 &= \{(0,0),(0,1),(0,2)\};\\
A_3 &= \{(0,0),(0,1),(0.25,0)\};\\
A_4 &= \{(0,0),(0,1),(0.25,0),(0,2),(0.5,0)\}.
\end{aligned}
$$

Each of these dilating elements determine the types of variations that can be detected. A_1 and A_2 both only deal with variations in pitch. A_1 allows us to detect variations of one semitone and A_2 of both one semitone and one whole tone. A_3 and A_4 deal with both pitch and rhythm. A_3 can detect variations of one semitone and variations of one quarter-note, while A_4 can also detect variations of one whole tone and of one half-note. While using a structuring

Figure 4.1: The dilating elements used for the algorithm.

element with more points allows us to discover more occurrences of a pattern, it is also more susceptible to finding redundancies and to overshoot.

Let us first look at the Fugue n°2 (BWV 847), where the subject contains 20 notes and appears 8 times [20] . The method of the authors to find the subject of the fugue provided the same result.

Using regular erosion to find MTEC conjugate pairs gives us two MTEC conjugate pairs: one gives us a pattern of 19 points with five occurrences and the other a pattern of 20 points with four occurrences, neither of which is satisfying. Using the OVE with respect to an MTEC pair on any of those MTEC conjugate pairs then gives us a new pair corresponding to a pattern with 20 points, and either six occurrences in the case where the dilating element is A_1 or A_3, or eight occurrences in the case where the dilating element is A_2 or A_4. Figure 4.2 illustrates this.

Thus, using A_2 or A_4 as the structuring element allows us to find the same result as M. Giraud, R. Groult, and F. Levé [20], which was the expected result. Two elements of interpretation can be gleaned from this fact. Firstly, our approach seems indeed capable of improving the results obtained from looking at MTEC conjugate pairs, even when using A_1 or A_3. Secondly, the fact that using A_2 or A_4 provided better results than A_1 or A_3 suggests that variations of a whole-note need to be taken into account in order to get a satisfying read on Bach's fugues.

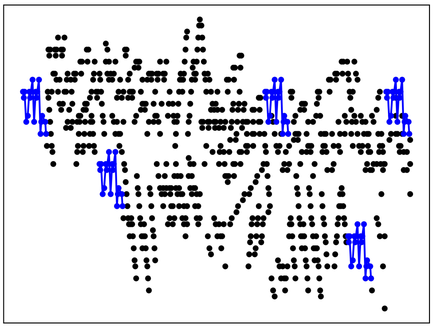

(a) In the Fugue n°2, we find an MTEC conjugate pair giving us a pattern with 19 points and 5 occurrences.

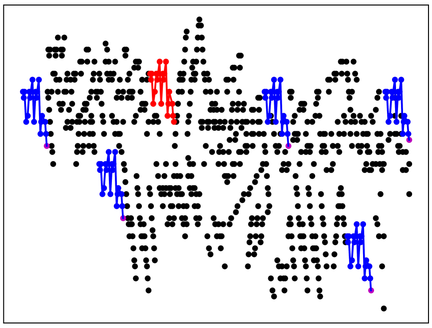

(b) Using the OVE with respect to the MTEC conjugate pair with dilating element A_1 or A_3 completes the pattern to 20 points and finds one more occurrence.

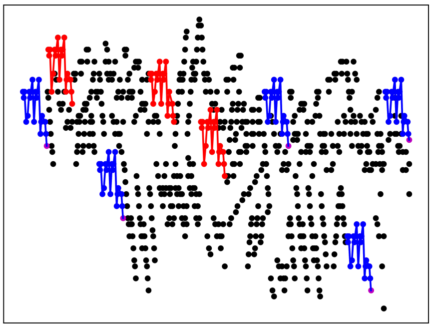

(c) Using the OVE with respect to the MTEC conjugate pair with dilating element A_2 or A_4 completes the pattern to 20 points and finds three more occurrences.

Figure 4.2: Completing an MTEC conjugate pair in Fugue n°2.

For another example where our approach leads to similar results than M. Giraud, R. Groult, and F. Levé, we can look at the Fugue n°1. However, as we will see, our approach allows us to glean some other information concerning the subject and its occurrences. The subject of this fugue contains 14 notes and appears 23 times [20]. Using their methods, the authors of the article found a subject of 14 notes that appeared 21 times.

With our method, using A_1, A_2 or A_3 we find five extended pattern-onset couples, and using A_4 we find six of them. The results are summarized in Table 4.1.

couple n°		A_1	A_2	A_3	A_4
1	N° of points	7	7	7	7
	N° of occurrences	23	23	23	25
2	N° of points				8
	N° of occurrences				24
3	N° of points	10	10	11	11
	N° of occurrences	22	22	22	23
4	N° of points	14	14	14	14
	N° of occurrences	21	21	21	22
5	N° of points	15	15	15	15
	N° of occurrences	6	6	9	22
6	N° of points	22	25	23	30
	N° of occurrences	3	3	3	3

Table 4.1: Results of the first algorithm on the Fugue n°1 for different dilating elements

Let us discuss the result of the A_1 column. The third element of the column gives us the same result as M. Giraud, R. Groult, and F. Levé in their article: a pattern with 14 points and 21 occurrences. However, looking at the first element gives us the expected number of occurrences, which was 23. Thus, our approach gives us some other information than that in the article of reference: it also gives us the number of certain partial occurrences of the pattern. Moreover, with the fourth and fifth element of the column, we also get occurrences of patterns that are greater than the subject and contain also parts of the counter subject. Using A_2 and A_3 gives way to similar results, with a few differences for the couples n°3, n°5 and n°6, while using A_4 gives us one more couple and generally makes us find either patterns that are greater or that have more occurrences. We illustrate the use of our approach for the couple n°1 and the couple n°4 in Figure 4.3 and Figure 4.4.

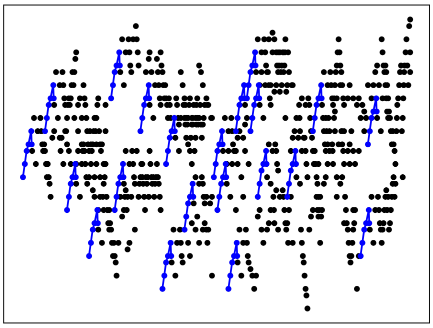

(a) In the Fugue n°1, we find an MTEC conjugate pair giving us a pattern with 6 points and 21 occurrences.

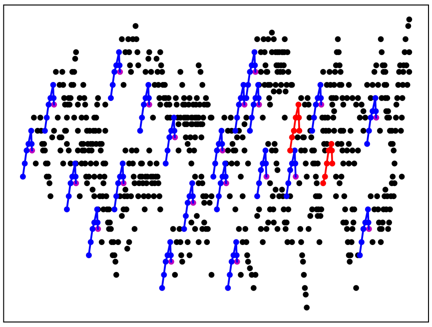

(b) Using the OVE with respect to the MTEC conjugate pair with dilating element A_1 or A_3 completes the pattern to 7 points and 23 occurrence.

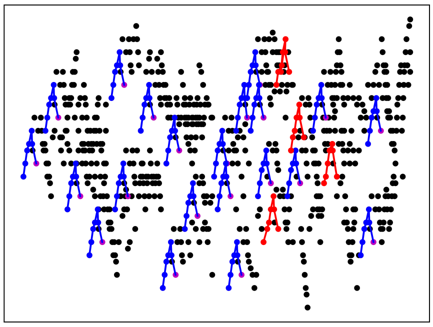

(c) Using the OVE with respect to the MTEC conjugate pair with dilating element A_2 or A_4 completes the pattern to 20 points 25 occurrences.

Figure 4.3: Completing an MTEC conjugate pair in Fugue n°1.

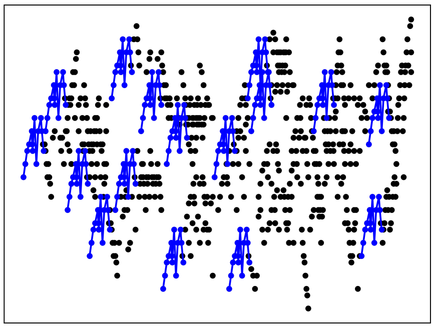

(a) In the Fugue n°1, we find an MTEC conjugate pair giving us a pattern with 13 points and 16 occurrences.

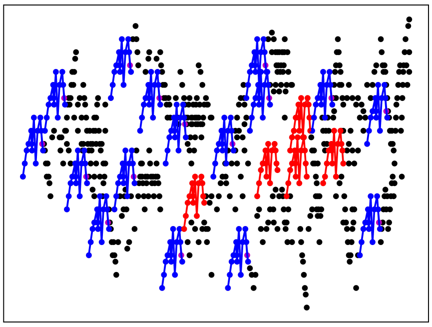

(b) Using the OVE with respect to the MTEC conjugate pair with dilating element A_1, A_2 or A_3 completes the pattern to 7 points and finds 21 occurrence.

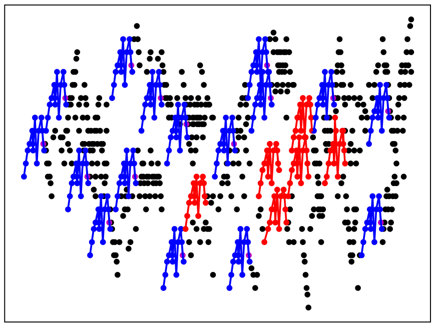

(c) Using the OVE with respect to the MTEC conjugate pair with dilating element A_4 completes the pattern to 14 points and finds 22 occurrences.

Figure 4.4: Completing an MTEC conjugate pair in Fugue n°1.

For a case where our approach proved to be an improvement of the results obtained by M. Giraud, R. Groult, and F. Levé, let us take a look at the Fugue n°5, where the subject contains 13 notes and appears 11 times [20]. It is a case where the authors of the article seem to have underestimated the number of notes in the subject and overshot the number of occurrences, as they have found a subject with 9 notes and 35 occurrences.

For every dilating element that we consider, we have found four extended pattern-onsets pairs. The results are summarized in Table 4.2.

In each case, the fourth pattern-onset pair gives us the result we were expecting: a pattern with 13 notes and 11 occurrences. But we should also note that the third pair gives us a pattern with 9 notes and that the second pair gives us a pattern with 35 occurrences (or 36 in the case where A_4 is used),

couple n°		A_1	A_2	A_3	A_4
1	N° of points	5	5	5	5
	N° of occurrences	37	37	37	56
2	N° of points	8	8	8	8
	N° of occurrences	35	35	35	36
3	N° of points	9	9	9	9
	N° of occurrences	21	21	23	23
4	N° of points	13	13	13	13
	N° of occurrences	11	11	11	11

Table 4.2: Results of the first algorithm on the Fugue n°5 for different dilating elements

which were the results found by M. Giraud, R. Groult, and F. Levé. Thus, our approach not only gives us the expected number of notes and number of occurrences, but also explains why M. Giraud, R. Groult, and F. Levé obtained the results they got. The Fugue n°5 is a piece where the subjects appears frequently in a truncated form. We illustrate our findings in Figure 4.5 and Figure 4.6, with our approach applied respectively to the second and fourth pair.

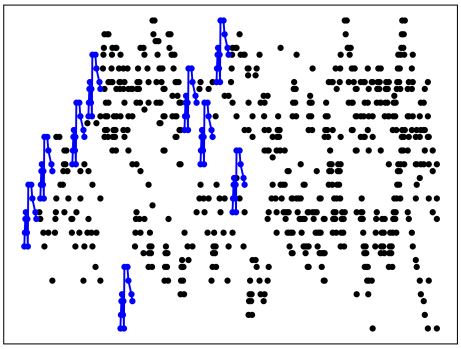

(a) In the Fugue n°5, we find an MTEC conjugate pair giving us a pattern with 13 points and 9 occurrences.

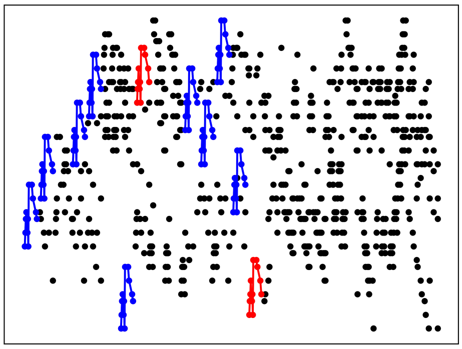

(b) Using the OVE with respect to the MTEC conjugate pair with dilating element any of our dilating elements keeps the pattern at 13 points and finds 11 occurrence.

Figure 4.5: Completing an MTEC conjugate pair in Fugue n°5.

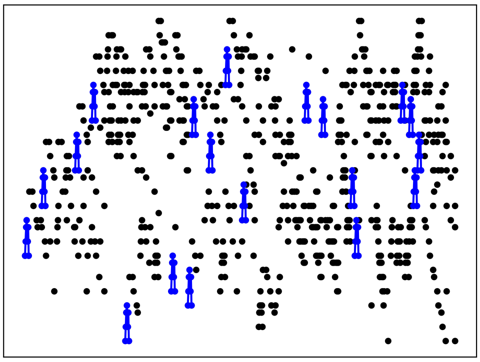

(a) In the Fugue n°5, we find an MTEC conjugate pair giving us a pattern with 8 points and 19 occurrences.

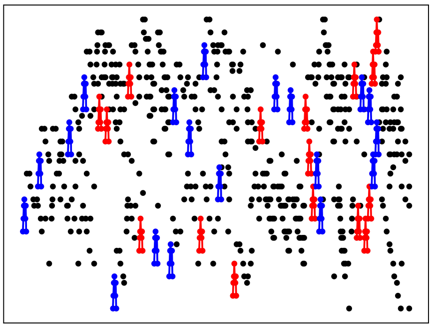

(b) Using the OVE with respect to the MTEC conjugate pair with dilating element A_1, A_2 or A_3 keeps the pattern at 8 points and finds 35 occurrence.

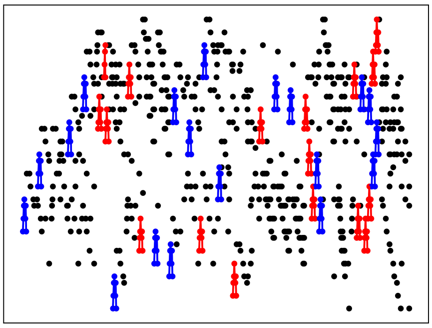

(c) Using the OVE with respect to the MTEC conjugate pair with dilating element A_4 keeps the pattern at 8 points and finds 36 occurrences.

Figure 4.6: Completing an MTEC conjugate pair in Fugue n°5.

One of the questions we need to ask ourselves is: which dilating element should be privileged when using our algorithm. Unfortunately, there does not seem to be an absolute answer to this question, which is the main limitation of our approach. To illustrate this, we can consider the Fugues n°9 and n°14. The Fugue n°9 has a subject with 6 notes and 12 occurrences [20], while the Fugue n°14 has a subject with 18 notes and 7 occurrences [20]. In both cases, M. Giraud, R. Groult, and F. Levé did not find the same result: they found a subject with 18 notes and 10 occurrences for the Fugue n°9 and a subject with 18 notes and 6 occurrences for the Fugue n°14.

In both cases, our approach enables us to find a pattern-onset pair with the expected values, but for only one dilating element. For the Fugue n°9, we need to use A_1, as the other dilating elements overshoot the number of occurrences. For the Fugue n°14 the correct result is obtained with A_4, as the other dilating elements find either not enough points in the pattern or not enough occurrences.

There thus does not seem to be a "size-fits-all" dilating element for our purposes. However, the dilating elements A_2 and A_4 seem to generally give way to more interesting results, as they take into account variations of one whole tone, which are not uncommon in Bach. A_4 does have a tendency to overshoot, but can also sometimes discover more relevant points than A_2.

Let us now discuss the limitations of our approach. We will first discuss the case of the Fugue n°16. According to our article of reference [20], the subject should contain 11 notes and should appear 16 times, and the authors found a subject with 11 notes and that appears 14 times, underestimating the number of occurrences.

Using our algorithm with A_4 gives us four pattern-onsets pairs, one of which gives us a pattern with 11 notes and 16 occurrences. One may think that this is a good thing, as we have found the expected number of occurrences. However, upon closer inspection this result comes from the fact that an occurrence has been counted several times. This is illustrated in Figure 4.7. What this means, is that the method we use to avoid counting redundancies is not full proof, and that they are cases where the results "lie". In this case, we thought we had improved the results of M. Giraud, R. Groult, and F. Levé, but in fact we only have found the same result.

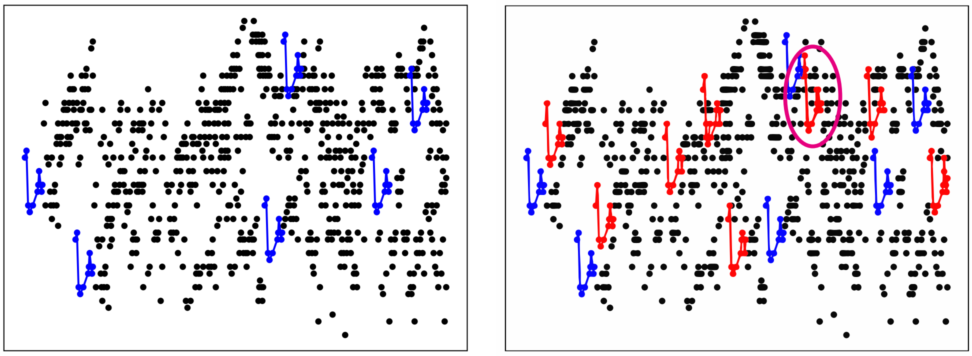

(a) In the Fugue n°16, we find an MTEC conjugate pair giving us a pattern with 11 points and 6 occurrences.

(b) Using our method with A_4 keeps the pattern at 11 points and finds 16 occurrence. The circled pattern is counted 3 times.

Figure 4.7: Completing an MTEC conjugate pair in Fugue n°16.

Finally, let us end this section with a case where our approach gives us significantly worse results than that of M. Giraud, R. Groult, and F. Levé. Let us consider the case of the Fugue n°21. The subject of this fugue has 38 notes and 8 occurrences. This is also the result discovered by the authors of the articles.

Using our approach, for each dilating element we find only one pattern-onsets pair. Using A_1, we find a pattern of 112 notes with only 2 occurrences. Using the other dilating elements, the pattern found is even greater, and still has 2 occurrences. In this case, we widely overshot the number of notes in the pattern. The reason is because the MTEC conjugate pair we attempt to extend already has a pattern with two many notes, i.e. the subject ended up mixed with the counter-subject giving us a subject that is "too long". This is illustrated in Figure 4.8. M. Giraud, R. Groult, and F. Levé bypassed this problem by fixing an upper limit to the number of notes the subject could have. Thus, this example illustrates a limitation of our approach: since we need to start from MTEC conjugate pairs in order to find patterns with variations, we cannot find patterns that are not extension of the patterns of the MTEC pairs. The next algorithm, using the PVE, does not have this problem.

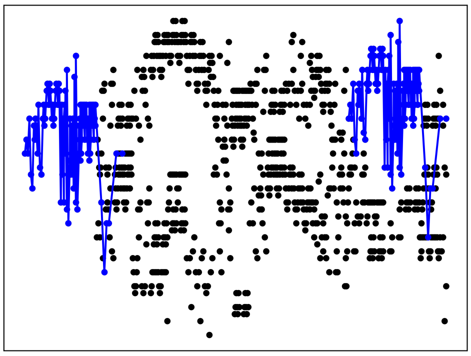

(a) In the Fugue n°21, we find an MTEC conjugate pair giving us a pattern with 109 points and 2 occurrences.

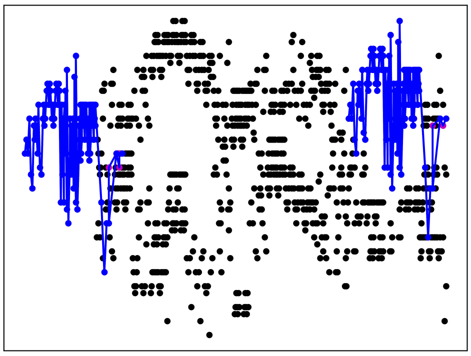

(b) Using our method with A_1 completes the pattern at 112 points and finds 2 occurrence.

Figure 4.8: Completing an MTEC conjugate pair in Fugue n°21

4.2 Using the PVE

4.2.1 Description

This second algorithm is based on the PVE. Since we do not need to start from MTEC conjugate pairs in order to use this operation, the process is different from the previous algorithm. Instead of applying the morphological erosion and then the variational operation, we will directly apply the PVE to the potential motifs.

This algorithm is to the first two steps of the previous one, with the PVE taking on the role of the regular erosion. The first step, which consists in finding pattern-onset pairs from the dataset is described by Algorithm 3.

Let X be the dataset we wish to study, and A the dilating element used to study variations. Starting with $n = 1$, we consider the three sets constituted from n consecutive notes, starting from the first, second and third note. For each one of these sets, we consider the PVE of X by the set, using the first note of the set as the pivot. This gives us the onsets of the set in X. If there is only one onset, then it means that the motif appears only once in the piece (with variations accounted for this time) and thus that it is not a repeating pattern. If this is the case for all three of the considered sets, we can stop the process, as any of the following sets we might consider will contain at least one of those and thus cannot be a repeating pattern. If there is more than one onset, then it means that it is a repeating pattern. We then compute the PVE of X by the onsets, choosing the origin of the space O_E as the pivot, in order to complete the pattern. Since the original set S is chosen such that $S \subset X$, choosing O_E as the pivot ensures that we can apply Theorem 3.2.1. and thus that we won't gain more information by applying the PVE with the same pivots.

We thus obtain a pair (P, O) with P a pattern and O its onsets. However, a pattern P can be found with different onsets depending on the pivot we have used. To ensure that a pattern is always associated to a unique set of onsets, we apply the PVE to X one last time, using the pattern P as the structuring element and the first element of P as the pivot, obtaining a new set of onsets O'. If (P, O') has not already be found, we add it to our list of pattern-onset pairs.

Algorithm 3 Finding patterns from consecutive notes with the PVE

$n \leftarrow 1$
$Pairs \leftarrow []$
$stop \leftarrow$ False
while not $stop$ **do**
  $onsets \leftarrow []$
  for $k \leftarrow 1$ to 3 **do**
    $S \leftarrow$ the n consecutive notes of X starting from the k-th note
    $s \leftarrow$ the first note of S
    $O \leftarrow \hat{\varepsilon}^{s}_{S,A}(X)$
    if $|O| > 1$ **then** insert O in $onsets$
    end if
  end for
  if $onsets$ is empty **then**
    $stop \leftarrow$ True
  else
    for $O \in onsets$ **do**
      $P \leftarrow \varepsilon^{O_E}_{O}(X)$
      $p \leftarrow$ the first note of P
      $O' \leftarrow \varepsilon^{p}_{P}(X)$
      if (P, O') not in $pairs$ **then**
        insert (P, O') in $pairs$
      end if
    end for
  end if
end while

The next and final steps consists in discriminating between patterns that are potentially musically interesting and those that are not. We can use the same set of heuristics as for the first algorithm: we consider the number of notes and the compactness of a pattern as significant parameters. We once again choose a minimum number of 4 and a compactness of 0.6 (with compactness being defined in the same way as before.

In the end, the algorithm gives us a list of pattern-onset pairs that give us a pattern and its occurrences in X with variations as defined by the dilating element A, similarly to the first algorithm.

The complexity necessary to compute the PVE of a set X by a structuring element P with dilating element A is $O(|A||P||X|ln(|A||P||X|))$, since it has equivalent complexity to doing the erosion of the dilated dataset by the pattern. Since we are doing at most three PVEs of X by every set consecutive notes in X up until a certain point, this gives us a worst case complexity of $O(k * n^3 ln(k * n)$, where $k = |A|$ and $n = |X|$.

4.2.2 Results

In this section we consider a few cases to illustrate some of the benefits and drawbacks of the approach and compare our results with those of M. Giraud, R. Groult and F. Levé [19], but also with those obtained with our first algorithm. We study the fugues from the Well-Tempered Clavier using 2 distinct dilating elements. Those are different from the ones used in the previous section because the PVE works better with symmetrical dilating elements. They are defined below and represented in Figure 4.9.

$$B_1 = \{(0,0),(0,1),(0,-1)\};$$
$$B_2 = \{(0,0),(0,1),(0,-1),(0,2),(0,-2)\};$$

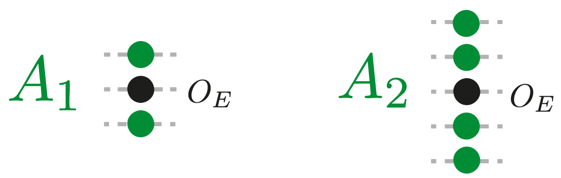

Figure 4.9: The dilating elements used for the algorithm.

The first thing one can remark concerning this algorithm is that we generally obtain more pattern-onsets couples than in our first approach. This is exemplified by the case of the Fugue n°2. The results obtained with both structuring elements are summarized in Table 4.3.

In the first algorithm, using A_1 or A_2 led to only one pattern-onset pair discovered: a pattern with 20 points and 6 occurrences in the case of A_1 and a pattern with 20 points and 8 occurrences in the case of A_2. In the approach using the PVE, we get 5 pattern-onset pairs using B_1, and 9 pattern-onset pairs using B_2. In Table 4.3, we can see that the fifth element found with B_1 is a pattern with 20 points and 6 occurrences and that the ninth element found with B_2 is a pattern with 20 points and 8 occurrences, confirming the results from the first algorithm. However, the other pairs also give us truncated instances of the subject.

This can be seen as a benefit of our approach. Discovering partial occurrences of a pattern can refine the analysis of the musical piece, especially in the context of Bach's fugues where the use of truncated instances of the subject is often deliberate, especially during the episodes. In Figure 4.10, we show two of the repeating patterns found using the PVE with B_2. The first one corresponds to the subject, and the second one to a truncated version of the subject with only 5 notes. In the second case, we can clearly see the episodes of the fugue appear.

couple n°		B_1	B_2
1	N° of points	5	5
	N° of occurrences	16	22
2	N° of points	5	5
	N° of occurrences	9	15
3	N° of points	10	5
	N° of occurrences	7	15
4	N° of points	6	6
	N° of occurrences	8	14
5	N° of points	20	6
	N° of occurrences	6	14
6	N° of points		11
	N° of occurrences		11
7	N° of points		8
	N° of occurrences		13
8	N° of points		8
	N° of occurrences		12
9	N° of points		20
	N° of occurrences		8

Table 4.3: Results of the second algorithm on the Fugue n°2 for different dilating elements.

However, having more pattern-onset pairs can also unnecessarily complexify the analysis of the fugue, and draw attention away from the important elements, which in our case are the subject and maybe a few of its truncated

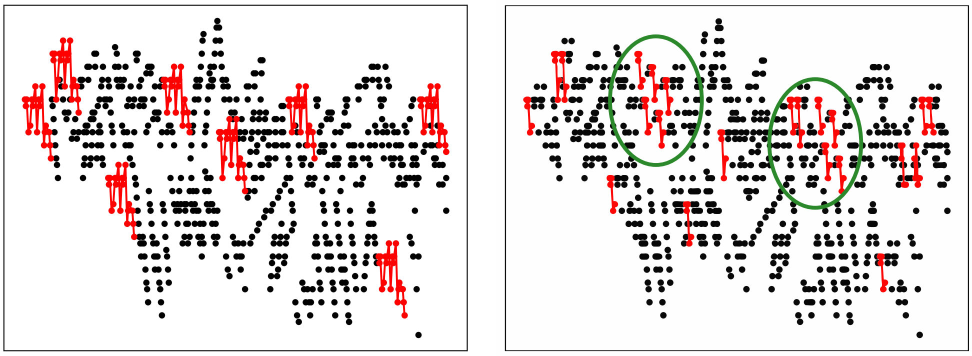

(a) We have found a subject with 20 notes and 8 occurrences.

(b) We have found that the subject truncated to the five first notes appears 22 times. The circled areas correspond to the episodes.

Figure 4.10: Patterns found using the PVE with B_2 in Fugue n°2.

instances. In order to select the most interesting pairs and avoid such a problem, we may want to add other criteria of selection in the future.

The example of the Fugue n°2 exemplifies pretty well the general behavior of the PVE-based algorithm in comparison to the OVE-based algorithm. Generally, the PVE-based algorithm finds more pattern-onsets pairs than the OVE-based algorithm, and the results found with the OVE-based algorithm are part of the results found with the PVE-based algorithm. When that is the case, this allows us to refine the analysis found with our first algorithm, but can also complexify it more than necessary. This behavior can notably also be observed for the Fugue n°1 and the Fugue n°5.

For an example where this approach leads to better results than the first algorithm, we may take a look at the Fugue n°21. Recall that when we used the OVE-based algorithm, whatever the dilating element used, we found only one pattern with more than one hundred points and only 2 occurrences, when we were expecting a pattern with 38 points and 8 occurrences. Using a PVE-based algorithm, we get the results that are summarized in Table 4.4.

One can thus remark that using B_2 as a dilating element allows us to discover a pattern with 38 points and 8 occurrences (the 10^{th} pair in the second column), which were the values expected for the subject of the fugue. Thus, the PVE-based algorithm can improve results in some cases where the OVE-based algorithm is unsatisfying. The pattern found is represented in Figure 4.11.

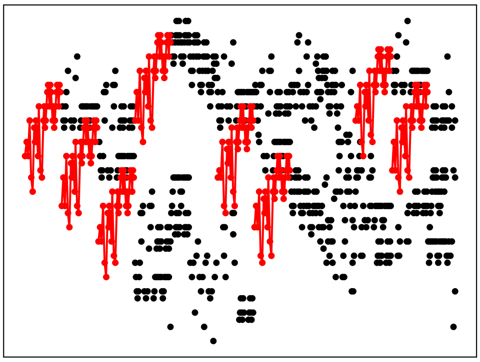

Figure 4.11: Using the PVE with dilating element B_2 in the Fugue n°21, we find a pattern with 38 notes and 8 occurrences.

couple n°		B_1	B_2
1	N° of points	5	5
	N° of occurrences	11	36
2	N° of points	5	5
	N° of occurrences	17	33
3	N° of points	5	5
	N° of occurrences	13	32
4	N° of points	38	8
	N° of occurrences	5	14
5	N° of points	8	6
	N° of occurrences	5	19
6	N° of points	39	6
	N° of occurrences	4	27
7	N° of points	112	7
	N° of occurrences	2	22
8	N° of points	90	12
	N° of occurrences	2	12
9	N° of points		8
	N° of occurrences		14
10	N° of points		38
	N° of occurrences		8
11	N° of points		97
	N° of occurrences		3
12	N° of points		109
	N° of occurrences		3
13	N° of points		121
	N° of occurrences		2

Table 4.4: Results of the second algorithm on the Fugue n°21 for different dilating elements.

This approach is however still susceptible to redundancies. We can see this if we consider the example of the Fugue n°16, which already posed a problem for the previous algorithm. Using B_2, we find a pattern with 11 notes and 16 occurrences, which were the expected values for the subject of the fugue. However, upon closer inspection, a pattern was counted twice, and we have actually found only 15 occurrences. Figure 4.12 illustrates this case.

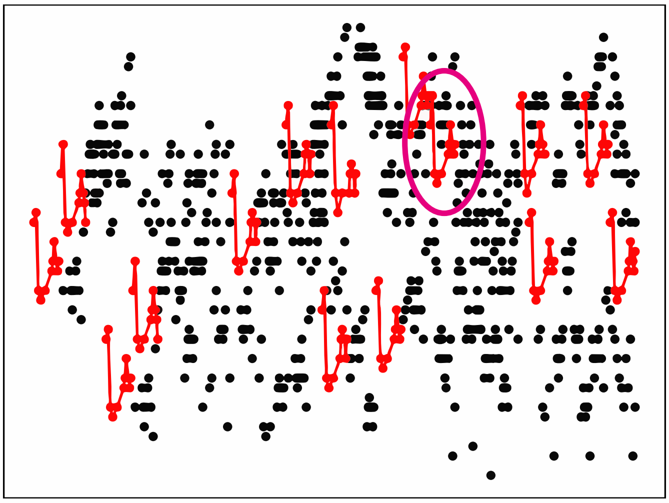

Figure 4.12: Using the PVE with dilating element B_2 in the Fugue n°16, we find a pattern with 11 notes and 16 occurrences. The circled occurrence was counted twice.

Chapter 5

Conclusion: Perspective and Future Works

In this thesis, we have defined three operations that allow us to discover patterns and their occurrences with variations in a music piece: the *Opening-less Variational Erosion* (OVE), the *Pivotal Variational Erosion* (PVE) and the *Intersectional Variational Erosion* (IVE). All of them are based on the use of a dilating element which defines the kind of variation we can discover. We then have designed two algorithms, the first one based on the OVE and the second one based on the PVE, that use those operations in order to find the subject of a fugue. In both cases, the algorithm returns a list of pairs representing a pattern and its onsets, finding not only the subject of the fugue but also relevant truncated instances of the subject. In the case of the OVE-based algorithm, we start by finding MTEC conjugate pairs from consecutive points at the beginning of the fugue and then extend them using the OVE with respect to the MTEC conjugate pair, while in the PVE-algorithm we apply the PVE directly to those consecutive points. Generally, the first algorithm gives us less pairs when compared to the second. Using the PVE-algorithm can thus refine, but also sometimes needlessly complexify the analysis. We get from both algorithms results similar to those obtained with the approach of M. Giraud, R. Groult, and F. Levé [20], provided we choose an appropriate dilating element.

One of the limitations of our approach is that it is tailored to a specific case, which is the problem of identifying the subject of a fugue. In order to find musically significant recurring patterns in any music piece, another approach needs to be considered. However, one of the more interesting aspects of our approach is that it fits very well in the paradigm of finding patterns through the discovery of MTEC pairs. Thus, as algorithms designed to discover MTEC pairs in a point-set are developed, they can be combined with the use of the OVE or the PVE in order to allow for the discovery of patterns with their variations.

Another limitation of our approach is that it only accounts for small variations of a pattern, be it in pitch or in rhythm. However, as pointed out during the introduction, those are not the only types of transformations one can encounter. Our method does not take into account transformations such as inversion or temporal dilation for example. Some of these transformations, when they can be represented by a bijection, could be discovered through the deformation of the dataset. It would then be plausible to use our approach to find transformed patterns in datasets that themselves have small variations in pitch or rhythm.

Finally, our approach could potentially be improved with the use of fuzzy mathematical morphology. Fuzzy mathematical morphology applies the principles of morphology on fuzzy sets, i.e. sets whose elements have degrees of membership. There are two ways one could use fuzzy logic to improve the operations we have defined. Firstly, a fuzzy dilating element could be used. This would allow to easily differentiate between perfect occurrences of a pattern and occurrences with variation. The other way we could use fuzziness would be with the pattern themselves. This could be a way to relate patterns that have some notes in common, e.g. truncated instances of a pattern.

Bibliography

[1] Heinrich Schenker. *Harmony*. Vol. 1. University of Chicago Press, 1954.

[2] Ian Bent and William Drabkin. *Analysis (New Grove Handbooks in Music)*. Macmillan, 1987.

[3] Fred Lerdahl and Ray S Jackendoff. *A Generative Theory of Tonal Music, reissue, with a new preface*. MIT press, 1996.

[4] Jean-Jacques Nattiez and Jonathan M Dunsby. "Fondements d'une sémiologie de la musique". In: *Perspectives of New Music* (1977), pp. 226–233.

[5] Nicolas Ruwet. *Langage, musique, poésie*. Éditions du seuil Paris, 1972.

[6] David Temperley. *The cognition of basic musical structures*. MIT press, 2004.

[7] Pierre-Yves Rolland. "Discovering patterns in musical sequences". In: *Journal of New Music Research* 28.4 (1999), pp. 334–350.

[8] David Meredith, Kjell Lemström, and Geraint A Wiggins. "Algorithms for discovering repeated patterns in multidimensional representations of polyphonic music". In: *Journal of New Music Research* 31.4 (2002), pp. 321–345.

[9] Isabelle Bloch, Henk Heijmans, and Christian Ronse. "Mathematical morphology". In: *Handbook of spatial logics*. Springer, 2007, pp. 857–944.

[10] Georges Matheron. *Random Sets and Integral Geometry*. J. Wiley and Sons, New York, 1982.

[11] Jean Serra. *Image analysis and mathematical morphology*. Academic Press, London, 1982.

[12] Jean Serra. *Mathematical Morphology. Volume II: theoretical advances*. Academic Press, London, 1988.

[13] Paul Lascabettes. “Mathematical models for the discovery of musical patterns, structures and for performances analysis”. PhD thesis. Sorbonne Université, 2023.

[14] Enguérand Tamagna. “Découverte de motifs musicaux avec variations en utilisant la morphologie mathématique”. MA thesis. Sorbonne Université, 2024.

[15] Tom Collins, Jeremy Thurlow, Robin Laney, Alistair Willis, and Paul Garthwaite. “A comparative evaluation of algorithms for discovering translational patterns in baroque keyboard works”. In: *Proceedings of the International Symposium on Music Information Retrieval* (2010).

[16] David Meredith. “COSIATEC and SIATECCompress: Pattern discovery by geometric compression”. In: *ISMIR-International Society for Music Information Retrieval Conference*. 2013.

[17] Tom Collins and David Meredith. “Maximal translational equivalence classes of musical patterns in point-set representations”. In: *Mathematics and Computation in Music: Proceedings of MCM 2013, Montreal, QC, Canada, June 12-14, 2013.,* Springer. 2013, pp. 88–99.

[18] David Meredith. “Understanding and compressing music with maximal transformable patterns”. In: *International Conference on Human-Computer Interaction*. Springer. 2023, pp. 309–325.

[19] Paul Lascabettes and Isabelle Bloch. “Discovering Repeated Patterns From the Onsets in a Multidimensional Representation of Music”. In: *International Conference on Discrete Geometry and Mathematical Morphology*. Springer. 2024, pp. 192–203.

[20] Mathieu Giraud, Richard Groult, and Florence Levé. “Detecting episodes with harmonic sequences for fugue analysis”. In: *ISMIR-International Society for Music Information Retrieval Conference*. 2012.