MUSICAL PATTERN DISCOVERY AT REGULAR TIME INTERVALS THROUGH MATHEMATICAL MORPHOLOGY: AN APPLICATION TO POLYMETRIC STRUCTURE ANALYSIS

Mathys DANIEL (mathys.daniel@lip6.fr) [1,2], **Paul LASCABETTES** (lascabettes@math.unistra.fr) [2], **Moreno ANDREATTA** (andreatta@math.unistra.fr) [2,3], and **Isabelle BLOCH** (isabelle.bloch@sorbonne-universite.fr) [1]

[1] **Sorbonne Université, CNRS**, *LIP6*, Paris, France
[2] **Université de Strasbourg, CNRS**, *IRMA, IRMIA++ & CREAA*, Strasbourg, France
[3] **Sorbonne Université, CNRS**, *STMS, IRCAM*, Paris, France

ABSTRACT

Musical pattern discovery is central in computational music analysis, as repeated structures reveal information about compositional intentions and listener perceptions. Traditional symbolic pattern discovery methods usually extract all repeating patterns in musical data but sometimes discover many patterns that are musically irrelevant. To overcome these limitations, we introduce in this paper a method based on mathematical morphology that specifically targets patterns repeated at exact regular time intervals. By analyzing the length of the extracted patterns, our approach provides insights into the rhythmic signature of a piece and enables automatic structural analysis. We apply this method to *Polygondwanaland* by King Gizzard & The Lizard Wizard to describe its rhythmic complexity and extensive use of polymetric structures. The proposed approach extracts fewer patterns than state-of-the-art algorithms while significantly reducing computation time. These results highlight both the analytical potential of the method for musicological study and its effectiveness as a computational tool for large-scale pattern discovery.

1. INTRODUCTION

This paper deals with the longstanding problem of automatic discovery of musical patterns, at the intersection of music theory, computational musicology, computer science and mathematics. Identifying the musical material that repeats within a piece is a fundamental prerequisite for the analysis of a musical piece [1], and more broadly, for extracting high-level information from musical works [2]. Moreover, repetition plays a central role in music: it intervenes in compositional processes [3], and has been shown to help creating listening experiences, and facilitate the memory process [2].

The problem of musical pattern discovery has been approached from multiple perspectives in both audio and symbolic domains. In the audio domain, A. Klapuri [4] provides a detailed overview of existing representations and discusses the limitations inherent to audio-based analysis for pattern discovery, notably the difficulty of obtaining stable and discriminative features for repeated-pattern extraction. In the symbolic domain, the survey by B. Janssen *et al.* [2] summarizes the state of the art as of 2019, reviewing the main algorithmic families used for structural discovery and pattern discovery. Despite decades of research, the task of identifying patterns that repeat in the data remains challenging due to the variety of musical styles, the presence of rhythmic and melodic variations, and the combinatorial explosion of possible candidate. Specific existing approaches rely on geometric methods, such as SIA (Structure Induction Algorithm) [5], its extension SIATEC (SIA Translational Equivalence Class) [5] or on compression-based methods, such as COSIATEC (COmpression SIATEC) [6]. These methods are highly effective at exhaustively extracting all possible repeating patterns from a musical piece but with no guarantee that these patterns are musically meaningful or perceptually relevant. Moreover, their computational complexity is cubic or quasi-cubic in number of points (i.e. notes). In this work, we build upon these methods by introducing constraints to guide pattern discovery toward a specific, musically grounded type of pattern, while improving computational efficiency.

A recent line of work [7–9] explores a reinterpretation of pattern discovery using mathematical morphology, a formalism traditionally used in image processing [10]. This framework connects morphological erosion with musical pattern discovery algorithms. It provides a geometric and algebraic viewpoint on the pattern structure, offering tools to better understand the nature of musical patterns, and to interpret the results of existing algorithms. The present work builds directly on this perspective, enhancing the existing theoretical framework with an algorithm which aims to extract patterns that repeat exactly at regular time intervals in symbolic music representations. This perspective allows us to design a conceptually simple yet computationally efficient algorithm, capable of capturing musically relevant patterns and provide information about the metric structure. It therefore provides new structural tools for computational music analysis.

Copyright: © 2026. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 Unported License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

2. DATA

2.1 Dataset Choice

To properly evaluate the validity of our algorithm in situations that are complex enough, two criteria guided our choice of musical material. First, we require music with original rhythmic structures and unconventional time signatures, which avoids the predominance of $\frac{3}{4}$, $\frac{6}{8}$, $\frac{12}{8}$ and $\frac{4}{4}$, commonly found in most popular music. Secondly, we need pieces exhibiting significant polymetric content. Polymeter is known to be difficult for human listeners to perceive and analyze [11], which reinforces the relevance of an automatic computational method capable of extracting interpretable rhythmic information about the multiple used time signatures. For these reasons, we selected *Polygondwanaland* by King Gizzard & the Lizard Wizard, a 10-track multi-instrumental album lasting 44 minutes. This album features uncommon metric constructions and extensive use of polymetric layering, making it an ideal test for structural analysis and rhythmic pattern discovery. Moreover, the masters of this album have been released to the public for free by the band under a *Creative Commons BY-ND* license, which notably permits its use for research purposes without restrictions. The band's open release policy ensures that we can freely rely on this material and on datasets derived from it for academic research. In particular, we used an online MIDI dataset of the album created by 8-bit Escapades, which provides a symbolic representation suitable for our framework and is freely available for scholarly use [1]. To date, no computational musicological study has been conducted on this album despite its metric complexity. The simultaneous presence of multiple polymetric layers makes the metric structure of individual tracks difficult to perceive, and a manual exhaustive annotation would be prohibitively time-consuming for a musicologist. Our method addresses both of these challenges: it provides an exhaustive and objective account of exact repeating patterns based on a clear and reproducible criterion, and it offers analytical insights that complement the manual analysis made by musicologists. This album therefore constitutes a challenging benchmark for our algorithm, and is a musically significant object of study in its own right, for which our approach offers new analytical insights.

2.2 Point-Set Representation

In our framework, each musical piece is represented as a two-dimensional point set $X \subseteq \mathbb{R}^2$, where each point encodes a note by an onset-pitch pair. The temporal axis corresponds to note onsets expressed in beats, while the vertical axis corresponds to pitch values. This representation follows standard symbolic-music conventions and provides an adequate theoretical model for the application of morphological operators. When processing a MIDI file with PrettyMIDI, the temporal onset values are provided in seconds. In order to improve rhythmic interpretability, these values are converted to beats using the following formula:

$$onset_{beats} = onset_{seconds} \times \frac{bpm}{60}. \quad (1)$$

The resulting beat onsets are then quantized with a step size of $\frac{1}{8\times3}$. This choice allows ternary rhythms (factor $\frac{1}{3}$) to be represented while preserving a fine temporal resolution equivalent to a sixteenth note (factor $\frac{1}{8}$). Although this quantization step is chosen heuristically, it remains consistent with observations by V. Iyer *et al.*, who report that the Western-music tatum typically lies around $\frac{1}{16}$ or $\frac{1}{24}$ [12].

This quantized two-dimensional representation constitutes the input for all subsequent morphological operations, enabling exact matching of musical patterns repeating at regular time intervals without tolerance thresholds. This representation applies to melodic instruments and to percussion, as each percussion hit is encoded as a distinct MIDI pitch value corresponding to a specific drum sound.

2.3 Filtering Irrelevant Tracks

To ensure that our analysis focuses on musically meaningful rhythmic structures, we remove several tracks from the dataset that are not relevant for our study. Specifically, we filter out tracks containing only a few isolated notes or long sustained notes, as these do not contribute to the detection of repeated rhythmic patterns. The following tracks are therefore manually excluded from the dataset:

- *Innercell* – Keys II
- *Polygondwanaland* – Keys I & Intro Keys
- *Deserted Dunes Welcomes Weary Feet* – Bass Key, Flute/Keys II, Keys IV & Flute
- *Searching* – Keys I & Keys II
- *The Fourth Colour* – Keys I

After applying these exclusions, approximately 100 tracks remain in the dataset for analysis, with many polymetric structures (up to five layers).

3. MORPHOLOGICAL ALGORITHM FOR DISCOVERING PATTERNS REPEATING AT REGULAR TIME INTERVALS

3.1 Erosion as a Basic Operation to Extract Pattern Positions

We start by reviewing the definitions and properties of morphological erosion that are useful in this article. Let $X \subseteq \mathbb{R}^2$ be a set of points and $S \subseteq \mathbb{R}^2$ be a structuring element. The erosion of X by S is defined as [10]:

$$\varepsilon_S(X) = \{x \in \mathbb{R}^2 \mid S_x \subseteq X\}, \quad (2)$$

where $S_x = \{s + x \mid s \in S\}$. In other words, the erosion extracts all positions where the pattern S occurs in X. An example of an erosion applied to a set of points is illustrated in Figure 1. Here, a structuring element S of four points repeats three times in X, so $\varepsilon_S(X)$ is composed of three points (in red in the figure).

[1] Dataset available at: https://8-bitescapades.bandcamp.com/album/polygondwanaland

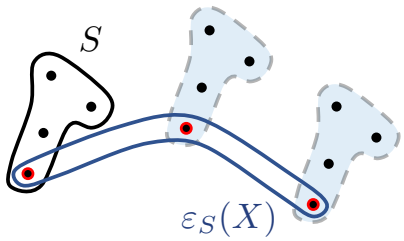


Figure 1. Application of morphological erosion applied to a set of points.

A fundamental property of erosion is that it *deflates the object* [10]. This means that erosion strictly removes points from X and never creates new ones, if and only if the origin of space belongs to S. Under this condition, erosion satisfies

$$\varepsilon_S(X) \subseteq X. \tag{3}$$

This property makes erosion particularly well suited for pattern extraction, as it guarantees that only points belonging to complete pattern occurrences are retained, thereby filtering out incomplete or irrelevant matches. Moreover, as the structuring element grows in size, this inclusion constraint becomes stricter, resulting in a more selective erosion ($S \subseteq S' \Rightarrow \varepsilon_{S'}(X) \subseteq \varepsilon_S(X)$) that identifies only the most significant patterns. This highlights the importance of choosing an appropriate structuring element, as its characteristics directly influence the nature of the subsets of points being extracted with an erosion.

3.2 Algorithm Overview

The proposed pattern discovery algorithm is mainly based on the specificity to adapt the structuring element of morphological erosion. In particular, it iterates over multiple structuring element lengths to detect patterns at different scales. Moreover, these patterns are further processed by `windowing` and `refine` operations. The general structure of the algorithm is presented in Algorithm 1, where each step as well as the notations are detailed in the following subsections [2] .

Algorithm 1 Morphological Pattern Detection Structure

1: **for** $L \in [2, 2.25, \ldots, L_{max}]$ **do**
2: $O_L \leftarrow \{(0,0),(L,0),(2L,0),(3L,0)\}$
3: $P \leftarrow \varepsilon_{O_L}(X)$ ▷ *Erosion with structuring element*
4: $P_{windowed} \leftarrow$ `windowing`(P, L)
5: $\{(P_j, N_j, L)\} \leftarrow$ `refine`$(P_{windowed}, P, L)$
6: **end for**
7: **return** Set of triplets (P_j, N_j, L)

Note that the `refine` operation may output a variable number of triplets, indexed by j, for each value of L, and that L may have decimal values.

[2] The code is available at `https://github.com/MathysDaniel/RegularTimeInterval_Patterns_Discovery`

3.3 Structuring Elements Choice

In this article, our goal is to detect musical patterns that repeat at regular temporal intervals. Following the morphological approach introduced by P. Lascabettes [7], a key idea is to use a temporally regular structuring element acting directly on onset positions. Given a temporal spacing L, we define the structuring element:

$$O_L = \{(0,0),(L,0),(2L,0),\ldots,(NL,0)\}, \tag{4}$$

which represents $N+1$ onsets separated by a fixed duration L. This construction directly encodes the hypothesis that a pattern repeats periodically.

A morphological erosion using O_L as structuring element then extracts the set of onsets at which a pattern P appears repeatedly with period L:

$$P = \varepsilon_{O_L}(X). \tag{5}$$

An example of a pattern P repeated four times at regular intervals L is illustrated in Figure 2. Therefore, by defining $O_L = \{(0,0),(L,0),(2L,0),(3L,0)\}$, we can obtain P from an erosion by O_L.

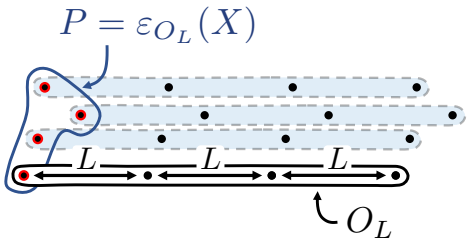


Figure 2. Set X (dots) and its morphological erosion by $O_L = \{(0,0),(L,0),(2L,0),(3L,0)\}$ which allows us to find a pattern P that repeats four times at regular intervals.

Additionally, the spacing L has a direct musical interpretation: it corresponds to the numerator of the time signature underlying the pattern (e.g., $L = 7$ suggests a $\frac{7}{4}$ or $\frac{7}{8}$ time signature), as mentioned in [7]. In this study, we fix $N = 3$, meaning that all structuring elements encode four repetitions, as illustrated in Figure 2. This musically grounded choice filters out patterns with too few occurrences to be analytically meaningful.

3.4 Subpattern Separation and Window Optimization

After computing the erosion $P = \varepsilon_{O_L}(X)$, the resulting point set may span up to the entire temporal duration of the piece, and may therefore contain several distinct subpatterns repeating with the same temporal spacing L. To separate them, we apply a window-based partitioning (Algorithm 2): the onset axis is divided into consecutive windows of length L, and each window containing at least one point of P is retained as a candidate subpattern. This procedure isolates subsets corresponding to individual occurrences of a repeated pattern.

Once the windows are determined, a refinement step (Algorithm 3) eliminates redundancies by exploiting the regularity of the extracted repetition structure. Since exact repetitions at interval L have been isolated through erosion,

Algorithm 2 Subsets Windowing (`windowing`)

Require: Point set P, window length L

Initialize windows $\leftarrow \emptyset$
while there remain points in P **do**
  Let t be the smallest onset in P
  Extract window $W \leftarrow \{p \in P \mid t \leq p_0 < t + L\}$
  if $W \neq \emptyset$ **then**
    Add W to windows
  end if
  Remove from P all points p with $p_0 < t + L$
end while
return windows

detecting additional occurrences consists in verifying, for each candidate window, whether all its points are included in another window translated in time by multiples of L.

For each pair of candidate windows P_j and P_k, P_j is translated to align with P_k, yielding $P_j^{(m)}$, and N_j denotes the number of occurrences of P_j, initialized to 4 by construction of the structuring element. If $P_j^{(m)} \subseteq P_k$, the pattern P_j is considered to occur once more and N_j is incremented. If additionally $P_k \subseteq P_j^{(m)}$, the two windows are equal up to a translation, an exact repetition has been found, and P_k is removed from the candidate list. A candidate subpattern P_k is considered only if its onset follows immediately (within tolerance) a previously validated occurrence at a distance of one window length L, ensuring that repetitions are detected as consecutive and uninterrupted. This approach is computationally efficient: a single erosion followed by inclusion tests suffices to identify all maximal subpatterns, avoiding the need for multiple erosion operations with varying structuring elements. The process yields a set of subpatterns that are both maximally informative and musically interpretable, each associated with its exact number of occurrences.

Figure 3 illustrates these two steps (windowing and refine) on a set of points. In this case, the algorithm returns a first pattern that repeats four times and a second pattern that repeats six times.

3.5 Compactness

The erosion and windowing steps generally produce several candidate patterns for a given spacing L, many of which may partially overlap or correspond to redundant descriptions of the same underlying musical structure. To resolve these ambiguities and ensure that each point of the dataset is assigned to at most one representative pattern, we introduce a *compactness* criterion, inspired by COSIATEC [6], that favors patterns offering the most informative and efficient description. Given a pattern P detected with N occurrences and temporal spacing L, we define the compactness as:

$$\mathrm{Comp}(P, N) = \frac{N|P|}{|P| + 2 + \ell}, \tag{6}$$

where $|P|$ denotes the number of points belonging to a single occurrence of the pattern, N is the number of repetitions of the pattern, the constant term $+2$ accounts for the two parameters needed to describe the structuring element (namely N and L), and ℓ counts the number of points of the original dataset X whose temporal coordinate lies within the temporal support of the pattern and its translations, but which do not belong to the pattern or any of its translations.

The compactness favors patterns that are both frequent (large N) and internally consistent (small ℓ), while also avoiding unnecessarily large or redundant descriptions. During the final selection stage, when two or more candidate patterns overlap or compete for the same points of X,

Algorithm 3 Repetition Maximization (`refine`)

Require: Window list P_{list}, original set X, length L

$P_{\text{to_test}} \leftarrow P_{\text{list}}$
$P_{\text{output}} \leftarrow \emptyset$
while $P_{\text{to_test}} \neq \emptyset$ **do**
  Select a window $P_j \in P_{\text{to_test}}$
  $N_j \leftarrow 4$
  $t_j \leftarrow \min\{t \mid (t, p) \in P_j\}$
  $R \leftarrow \{t_j\}$ ▷ *Reference onsets*
  for all $P_k \in P_{\text{list}} \setminus \{P_j\}$ **do**
    $t_k \leftarrow \min\{t : (t, p) \in P_k\}$
    if $\nexists r \in R$ such that $(t_k - r) = L$ **then**
      continue
    end if
    $m \leftarrow \lfloor (t_k - t_j)/L \rfloor$
    $P_j^{(m)} \leftarrow \{(t + mL, p) \mid (t, p) \in P_j\}$
    if $P_j^{(m) \subseteq P_k}$ **then**
      $N_j \leftarrow N_j + 1$
      Add t_k to R
      if $P_k \subseteq P_j^{(m)}$ **then**
        Remove P_k from $P_{\text{to_test}}$
      end if
    end if
  end for
  Add (P_j, N_j, L) to P_{output}
  Remove P_j from $P_{\text{to_test}}$
end while
return P_{output}

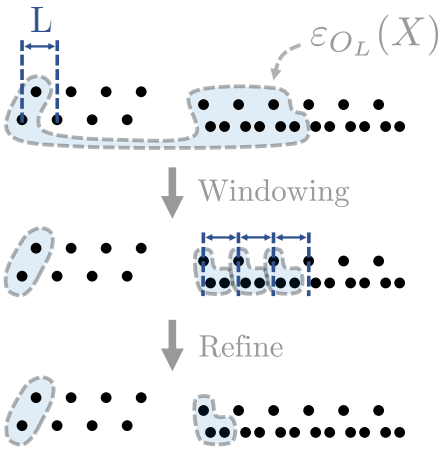


Figure 3. Illustration of the two additional steps in the algorithm: *windowing* (separating patterns) and *refine* (avoiding redundancies).

the algorithm retains only the one with the highest compactness score. This ensures that the final output consists of a set of coherent, non-overlapping, and musically meaningful patterns, each optimized to provide the most compact representation of the repeated structure it encodes. It significantly improves the interpretability of the results by favoring patterns offering the most informative and efficient description.

3.6 Computational Complexity Analysis

In this section, we analyze the runtime behavior of the proposed morphological algorithm, and we compare it experimentally with SIA. To this end, we measure the computational runtime of both algorithms on 8 tracks of the dataset. The SIA algorithm has been reimplemented in Python, and both algorithms are run on the same computer and with the same data, for the sake of comparison.

Figure 4 displays the computational time (in seconds) obtained for both algorithms independently, for various input sizes $n = |X|$. Simple regression models are fitted combining polynomial and $n \log n$ terms using least-squares estimation.

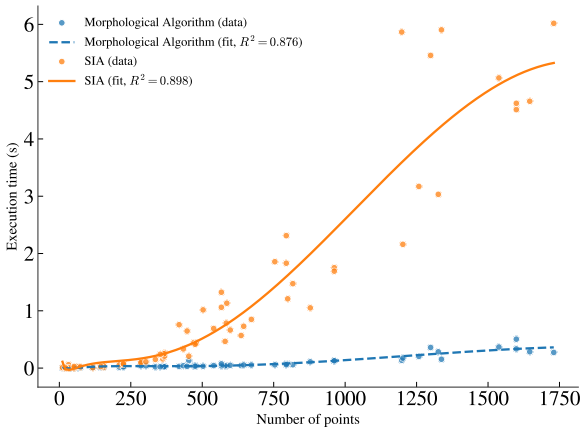


Figure 4. Comparison between the proposed morphological algorithm and SIA.

As shown in Figure 4, the proposed morphological algorithm is consistently faster than SIA across all tested input sizes. The observed speed-up ranges from approximately 1 to 45, with an average factor of 11.8 in favor of the proposed method. Moreover, the curvature of the fitted trends suggests that SIA follows a superlinear to near-quadratic growth, consistent with its known worst-case complexity in $\mathcal{O}(n^2 \log n)$ [5]. By contrast, the proposed algorithm exhibits a milder growth rate, which can be reasonably bounded between $\mathcal{O}(n)$ and $\mathcal{O}(n^2)$ in practice.

These observations are derived from experimental measurements on a specific corpus and are therefore valid only over the range of input sizes considered in our experiments; as such, they should not be interpreted as asymptotic complexity bounds, but rather as an indication of the practical runtime behavior observed in our study.

We now provide the worst case complexity of the proposed method. Let n be the number of points in the input dataset, and p the number of points that are stored after an erosion for a given value of L. The algorithm consists of three main operations per value of L.

1. Morphological erosion. The erosion scans all n points and therefore has cost $\mathcal{O}(n)$.

2. Windowing. The extraction of windows from this ordered list is linear: $\mathcal{O}(p)$.

3. Translation-based refinement. The refinement step compares each window against all other detected windows. If W is the number of windowed patterns, and each window i contains p_i points (with $\sum_{i=1}^{W} p_i = p$), the temporal contiguity constraint restricts comparisons to windows whose onset lies exactly one period L ahead of a previously validated occurrence, so each window P_j is effectively compared against only a constant number of candidates. The total cost of all inclusion tests is therefore $\mathcal{O}(p)$.

Compactness filtering. The compactness must be computed for each sub-pattern retained after the translation based refinement. This computation requires translations of the sub-pattern on the onsets of the corresponding structuring element (whose length equals the number of repetitions of the sub-pattern) to count the lost points. Since this operation is performed a limited number of times, its cost is negligible compared to $\mathcal{O}(p)$. The runtime measurements presented in Figure 4 correspond to the core algorithm (erosion, windowing, and refinement) without the compactness filtering step, which is applied separately for result interpretation. However, including this post-processing step would not change the asymptotic complexity analysis, as justified above.

Putting everything together, and assuming $\mathcal{L}$ possible structuring element lengths, we obtain the overall worst-case complexity:

$$\mathcal{O}\big(\mathcal{L}(n + p)\big)\,. \tag{7}$$

The experimental cost thus depends on two main factors: the number of points n in the musical piece and the lengths p of extracted patterns. In practice, the dominant factor is n, since p is generally much smaller than n as it corresponds to the length of patterns extracted from the full point set. The overall complexity may exhibit quadratic behavior when $\mathcal{L}$ approaches n, but tends to become linear for musical pieces with a large number of points or when the number of tested pattern lengths is optimally chosen. This reduction in computational cost compared to SIA stems from the fact that the proposed algorithm targets a specific, predefined type of pattern rather than performing an exhaustive search over all possible repeated structures. The goal is therefore not exhaustiveness, but the extraction of patterns that are a priori known to be musically meaningful for this type of analysis.

4. RESULTS

In this section, the main outcomes of the proposed algorithm are presented by applying it to the *Polygondwanaland* dataset described in Section 2. Our evaluation focuses on three complementary aspects: (i) the interpretability of the extracted rhythmic structures; (ii) the ability of the method to analyse polymetric passages that are challenging

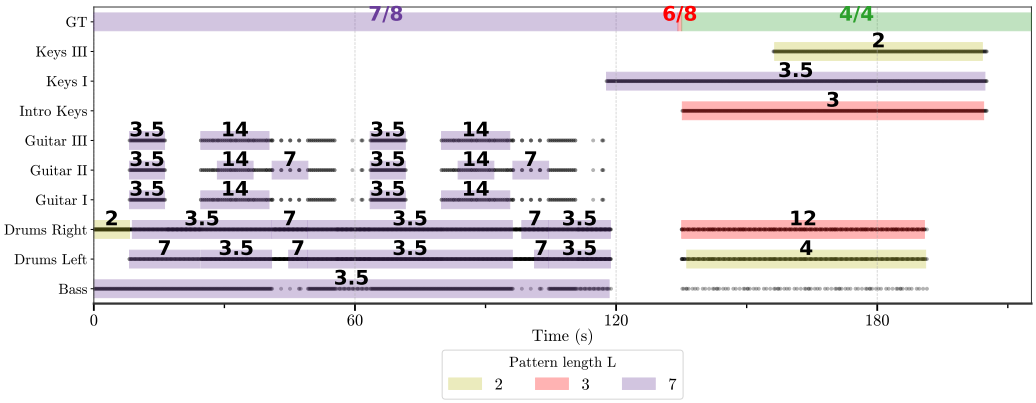


Figure 5. Structural analysis of *Deserted Dunes Welcomes Weary Feet* by King Gizzard & The Lizard Wizard.

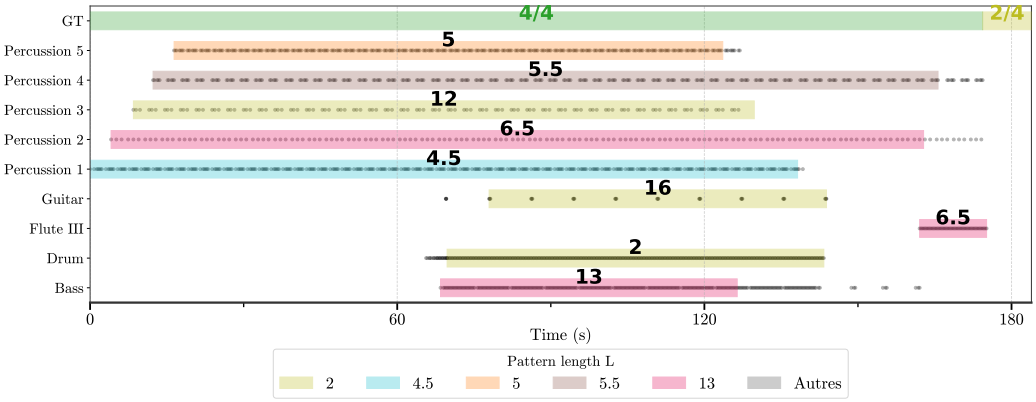


Figure 6. Structural analysis of *Searching* by King Gizzard & The Lizard Wizard.

to process manually; (iii) the reduction in the number of detected patterns compared to state-of-the-art approaches.

4.1 Pattern Length Structure

We illustrate the behaviour of the algorithm on three different tracks from the album: *Deserted Dunes Welcomes Weary Feet*, which feature homometric and polymetric structures; *Searching*, which contains highly pronounced polymetric interactions that are very difficult to perceive without stem separation; and *Innercell*, which presents an interesting polymeter structure. The analysis are made for $L_{max} = 20$.

Figure 5 shows the pattern lengths extracted for *Deserted Dunes Welcomes Weary Feet*. The algorithm clearly identifies its general structure: the first half is dominated by a strong septuple meter, easily observable in the extracted pattern lengths, while the second half reveals a polymetric superposition of triple, quadruple and septuple meters. However, it is important to note that the first part of the piece features what can be described as a 3-against-7 polymetric structure with the vocal part, which is not captured in this analysis but would appear if the voice stem were included in the dataset. The algorithm nonetheless captures the instrumental periodicities without requiring any strong modeling assumptions, providing clear insight into the metric organization of the piece.

Figure 6 displays the different pattern lengths found in *Searching*, a track that exemplifies high polymetric complexity. The analysis reveals a polymetric structure involving five distinct metric layers: 9, 10, 11, 12, and 13. Each instrumental track maintains its own time signature throughout the entire piece without deviation. This polymetric configuration is very difficult to perceive without isolating individual stems, which highlights the usefulness of analytical methods for understanding the construction of the piece.

Figure 7 presents the analysis of *Innercell*, which features a polymeter between triple and quintuple meters maintained throughout the track. This polymetric relationship is perceptible when listening to the interplay between guitars and drums, but our representation provides a clear structural view of the piece organization. The analysis shows that the guitars primarily operate in triple meter with a notable passage in quintuple meter. It is worth noting that in the introduction, the guitar part actually follows a pattern of seven measures in triple meter followed by one measure in quadruple meter. However, this single occurrence is below our minimum repetition threshold and is therefore not extracted. Nonetheless, the resulting representation remains musically relevant, as listener perception indeed retains the dominant triple meter, demonstrating that our approach aligns well with perceptual salience.

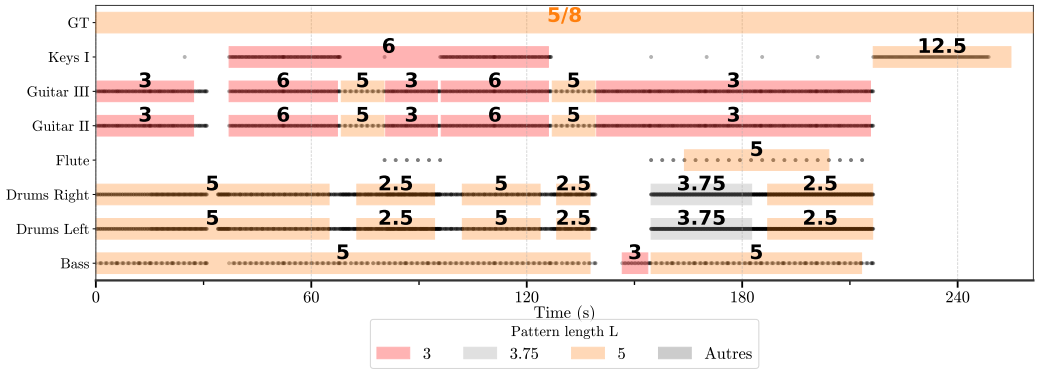


Figure 7. Structural analysis of *Innercell* by King Gizzard & The Lizard Wizard.

4.2 Pattern Length Distribution

Using the structural information about pattern lengths in musical pieces, it is possible to determine the proportion of polymetric and homometric sections in a given piece, as well as the pattern lengths used simultaneously, providing insight into the rhythmic signatures and their concurrent usage. For each musical piece, we analyze the entire temporal range to identify which pattern length(s) are active at each time instant.

Figure 8 (top) shows the distribution of combinations of pattern lengths for *Deserted Dunes Welcome Weary Feet*. A substantial portion of this piece clearly exhibits a homometric structure with pattern of length 7. It is observed that the septuple meter is always present throughout the entire piece and the polymetric structure alternates with pattern of lengths 2 and 3, either all together or separately. While this information was visible in the previous representation, it becomes even clearer in this visualization, with quantified proportions throughout the piece. This figure provides the necessary information to understand the metric distribution of the musical work.

Figure 8 (bottom) illustrates the distribution of combinations of pattern lengths for the piece *Innercell*. The analysis reveals that over half of the piece exhibits a 3-against-5 polymetric structure. The remainder consists of sections with pattern length 5 alone, pattern length 3 alone, or pattern length 3.75 (which represents one quarter of 15, corresponding to sections in quintuple meter where the most compact pattern, according to our chosen compactness definition, has length 3.75). The metric information is thus successfully extracted, providing a clear and readable view of the metric construction and its proportion of polymetric structures.

The figures for the other tracks on the album showing the structural analyses (as in Figure 7) and the distributions (as in Figure 8) are all available on GitHub. In particular, the analysis of *Searching* shows that 92% of the piece exhibits polymetric structures, with more than 50% of the track featuring the simultaneous presence of five pattern lengths (9, 10, 11, 12, 13).

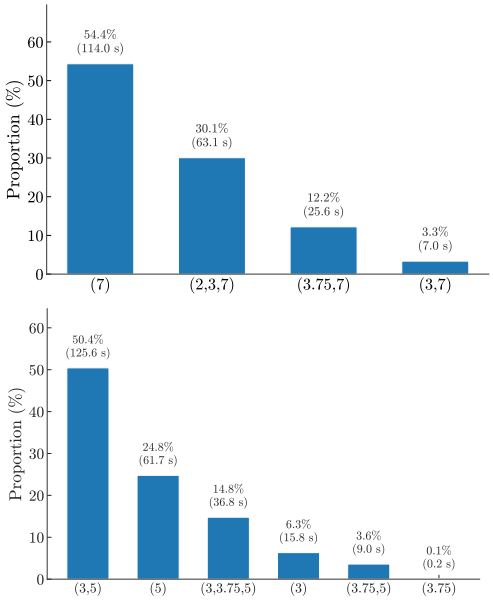


Figure 8. Distribution of simultaneous pattern lengths in *Deserted Dunes Welcome Weary Feet* (top) and in *Innercell* (bottom).

4.3 Amount of Patterns

A major advantage of our approach is its drastic reduction in the number of detected patterns, while still retaining musically meaningful structures. We compare the number of patterns produced by our algorithm to those obtained using SIA and SIATEC. Across the full *Polygondwanaland* album, our approach produces on average 252 times fewer patterns than SIA and 146 times fewer patterns than SIATEC, even before applying the compactness criterion. In typical cases, this corresponds to reduce the output from several tens of thousands of patterns down to a few hundred at most. This difference stems from the use of regular structuring elements, which target exact repetitions at regular time intervals, combined with the compactness criterion that further reduces redundant descriptions. This substantial reduction not only improves interpretability but also facilitates large-scale musicological studies where dense pattern outputs from classical methods become more difficult to inspect manually.

5. CONCLUSIONS

We have introduced a new morphological framework for symbolic musical pattern discovery based on erosions with temporally regular structuring elements, window-based subpattern separation, and a compactness-driven selection of representative patterns. The method offers a conceptually simple yet powerful alternative to classical geometric approaches.

Our experimental evaluation on the rhythmically complex album *Polygondwanaland* demonstrates several key advantages of the proposed approach:
(i) it uncovers musically meaningful periodic structures directly linked to underlying time signatures;
(ii) it drastically reduces the number of detected patterns, giving concise and interpretable rhythmic descriptors;
(iii) it achieves a reasonable computational complexity, making large-scale applications feasible;
(iv) it provides an exhaustive enumeration of repeating patterns that complements manual analysis made by musicologists.

These results highlight the relevance of mathematical morphology as a tool for symbolic music analysis, and confirm that regular-interval structuring elements provide a natural and musically grounded way to capture periodicity and metric organization.

However, it is important to note that our evaluation was conducted on a dataset particularly suited to periodic pattern discovery. Moreover, the dataset lacks vocal tracks, which contribute significantly to the polymetric perception in King Gizzard & the Lizard Wizard's music. Completing the dataset by incorporating vocal transcriptions would provide a more comprehensive representation of the rhythmic complexity.

Future work will focus on several key directions. We will extend the morphological framework and algorithms to handle patterns with controlled variations. Recent work has demonstrated both the feasibility of extracting approximate patterns [13] and the effectiveness of mathematical morphology as a tool for approximate pattern detection [14], suggesting promising directions for this extension. The robustness of pattern length detection could also be improved by incorporating contextual information, and therefore enhance the musical interpretability of the results. Another direction concerns the optimal selection of the algorithm's hyperparameters: the minimum number of repetitions N and the range of pattern lengths L to test. Finally, while no musicological study has been conducted on *Polygondwanaland*, our current evaluation remains primarily algorithmic. A collaborative analysis with music theorists would provide valuable insights into the musical significance of the detected patterns and could reveal additional structural relationships not captured by the computational approach alone. Applying the proposed method to larger and more diverse datasets would also be a natural next step to further assess its generalizability.

6. REFERENCES

[1] I. D. Bent and A. Pople, "Analysis," *Grove Music Online*, 2001.

[2] B. Janssen, W. B. de Haas, A. Volk, and P. van Kranenburg, "Finding repeated patterns in music: State of knowledge, challenges, perspectives," in *Int. Symposium on Computer Music Interdisciplinary Research*, vol. LNCS 8905, 2013, pp. 277–297.

[3] G. T. Jones, *Music Theory*. New York, Barnes & Noble Books, 1974.

[4] A. Klapuri, "Pattern induction and matching in music signals," in *Int. Symposium on Computer Music Interdisciplinary Research*, vol. LNCS 6684, 2010, pp. 188–204.

[5] D. Meredith, K. Lemström, and G. A. Wiggins, "Algorithms for discovering repeated patterns in multidimensional representations of polyphonic music," *Journal of New Music Research*, vol. 31, no. 4, pp. 321–345, 2002.

[6] D. Meredith, "COSIATEC and SIATECCOMPRESS: Pattern discovery by geometric compression," in *Int. Society for Music Information Retrieval Conference*, 2013, pp. 1–6.

[7] P. Lascabettes, "Mathematical models for the discovery of musical patterns, structures and for performances analysis," PhD Thesis, Sorbonne Université, 2023.

[8] P. Lascabettes and I. Bloch, "Discovering Repeated Patterns from the Onsets in a Multidimensional Representation of Music," in *Discrete Geometry and Mathematical Morphology*, S. Brunetti, A. Frosini, and S. Rinaldi, Eds., 2024, vol. LNCS 14605, pp. 192–203.

[9] ——, "Using mathematical morphology to discover repeated patterns in music," *Journal of Mathematics and Music*, pp. 1–26, 2025.

[10] J. Serra, *Image Analysis and Mathematical Morphology*. Academic Press, London, 1982.

[11] J. London, *Hearing in Time: Psychological Aspects of Musical Meter*. Oxford University Press, 2012.

[12] V. Iyer, J. Bilmes, M. Wright, and D. Wessel, "A novel representation for rhythmic structure," in *Int. Computer Music Conference*, 1997, pp. 97–100.

[13] D. Meredith, "Understanding and compressing music with maximal transformable patterns," in *Int. Conference on Human-Computer Interaction*. Springer, 2023, pp. 309–325.

[14] P. Lascabettes, N. Quaetaert, M. Daniel, M. Andreatta, and I. Bloch, "Applying morphological operations to subsets of points for the discovery of repeated musical patterns and their variations," *Journal of Mathematical Imaging and Vision*, vol. 68, no. 3, pp. 1–27, 2026.