Music Score Analysis with Process Mining (Extended Abstract)

István Koren

Chair of Process and Data Science, RWTH Aachen University, Aachen, Germany

Abstract

Process mining is applied to a wide variety of use cases, most typically for processes like order-to-cash and purchase-to-pay. Respective algorithms are able to discover bottlenecks, identify recurring patterns and deal with concept drifts. So far, use cases in cultural heritage scenarios are scarce. Music scores provide an ideal opportunity for process mining algorithms; they are structured into notes, measures, repetitions, parts, instruments, etc. In addition, there are compelling characteristics, like varying dynamics, transformations through modulations, and phase delays as in a fugue. In this demo article, we present a tool that is able to do a process-style analysis of music scores. Most notably, it transforms scores into an event log and performs basic process discovery. We see potential not only for music theorists but also as a pedagogical tool to illustrate process mining concepts and as a means to produce event logs for advancing process mining algorithms and visualizations.

Keywords

Process Mining, Music Analysis, Cultural Heritage

1. Introduction

Process mining is used in a wide variety of use cases for all kinds of industrial and societal processes [1]. Event logs are at the baseline; they contain sequences of events that describe the execution of a process. Taking event logs as starting point, process mining can be used to automatically discover, monitor, and improve processes. For example, it can discover how a process is actually being executed, find bottlenecks, or identify improvement opportunities. Process mining is a relatively new field, yet it is a valuable tool for any organization that wants to improve its business processes.

Use cases in the field of cultural heritage are rare so far. For instance, process mining techniques have been compared to the methodological tool *chaîne opératoire* common in archaeology [2]. We have not found any applications of process mining in music, despite the obvious parallels: Music pieces are usually denoted in a strongly formalized notation called scores. Especially in orchestral pieces, there is a lot of concurrency of instruments; throughout history, there have been a number of genres with similar patterns like repetitions and variations.

The tool and a demo video are available at https://istvank.eu/musicpm

ICPM 2022 Doctoral Consortium and Tool Demonstration Track

koren@pads.rwth-aachen.de (I. Koren)
https://pads.rwth-aachen.de (I. Koren)
0000-0003-1350-6732 (I. Koren)

We postulate the following research question: Can music scores be processed so that existing process mining tools can be used to analyze musical pieces?

In this demo paper, we present a proof-of-concept tool that is able to transform musical scores in the popular MusicXML notation to event logs. This enables new interpretations of musical art, for instance, by analyzing rhythm and repetitions in a process-oriented visualization. Specifically, we encode a piano piece of a recent number one hit single as event log. The visualized results are surprisingly clear. Therefore, we are optimistic that this type of analysis can help advance not only music theory but also process mining by making the large body of music scores available as event logs. Our tool can also be used as a pedagogical instrument to vividly explain business process management concepts.

2. Parallels to Process Analytics

Regarding music and data mining, there are not many related work applying standard data mining methods to music; a notable exception is the book Music Data Mining [3]. Baratè et al. present a musicological analysis with Petri nets as formal tools for studying concurrent, asynchronous, and parallel processes [4]. They focus on extracting groups from music scores as objects, such as episodes, themes, and rhythmic patterns, to visualize them as transition system. *Musical set theory* as a subdiscipline of music theory similarly organizes musical objects and describes their relationships to discover deep structures (e.g., [5]).

Especially in orchestral pieces, there is a lot of concurrency by instruments played simultaneously; throughout history, there have been many genres with similar patterns, like repetitions and variations. We postulate the following research question: Can music scores be processed so that existing process mining tools can be used to analyze musical pieces?

Chord: Group of multiples notes played together; stands for activity variants.

Dynamics: The volume of (a sequence of) notes; could represent the amount of resources consumed.

Modulation: Refers to a key change, if a melody is replayed on a different base note; corresponds to concept drifts.

Note: Basic unit of a score, besides pauses; can be considered the most fine-grained activity.

Repetition: Looped parts, often explicitly denoted variations; stands for re-executions.

Rhythm: A sequence of beats forming recognizable patterns; could represent the timing of events.

Tempo: How fast the piece is played, corresponds to the execution speed.

Typically, these features are combined in infinite possible variations. Throughout history, particular styles have formed, such as the fugue, where a musical theme is repeated in various pitches and through different instruments. It is trivial to see further parallels to process mining techniques such as frequent pattern mining and local process models.

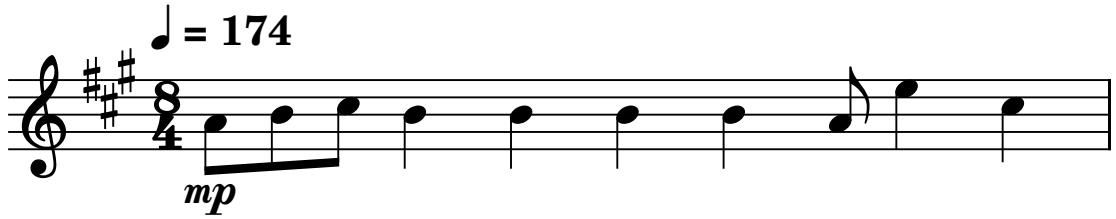

Figure 1: Characteristic first and last measure from "As It Was" by Harry Styles; this measure corresponds to an activity.

3. Mining Approach and Example

Our general approach is as follows. First, we parse a digital music score file. Currently, the entire piece of music is considered as one case. In a preprocessing step, we either select a single part (e.g., the left or right hand in a piano piece or a specific instrument in a classic orchestra piece) or merge all notes played in parallel. We then step through the score, following repetitions, and create one activity per unique measure, c.f. Figure 1. As the activity notion, we serialize the measure and reuse existing activities if the same musical pattern has already been discovered. Based on the annotated tempo, we calculate a timestamp for each measure.

Our implementation is based on a Python script using the music21 [6] and PM4Py [7] libraries. The library can parse a number of digital music score formats, including the popular MusicXML, an XML-based data format. We map the input of the measures as PM4Py Event objects. PM4Py is then used to export the log as XES file and display the resulting process model.

As an example piece, we chose a relatively simple piano arrangement of the worldwide number-one hit "As It Was" by Harry Styles as a MusicXML file. It is a suitable example, as the characteristic sequence shown in Figure 1 fits in one measure and is repeated several times. The resulting process model is shown in Figure 2 as a directly-follows graph (DFG). It is very easy to see that the melody known from Figure 1 (where the left hand has a break) forms the first and the last measure of the song. In addition, several repetitive patterns can be recognized, including a large one, as well as the repeated succession of the characteristic sequence $\{m5, m6, m7, m8\}$. On the web frontend, the activities in the process model can be clicked to replay the corresponding notes as audio.

4. Discussion and Outlook

In this paper, we positioned process mining as a technique to analyze music. This is possible since music scores share many characteristics with business processes. As a proof-of-concept, we presented a tool that transforms musical notation into event logs so that hundreds of process mining techniques can be applied directly. Therefore, we can answer the research question from Section 1; it is possible to repurpose process mining tools for music scores by performing a number of preprocessing steps. In the demo, we analyzed the 2022 hit "As It Was" by Harry Styles by transforming a piano arrangement as an event log and creating a DFG.

The current algorithm that turns music scores into process models follows a very simple scheme by taking measures as input and assigning them labels sequentially. As a next step,

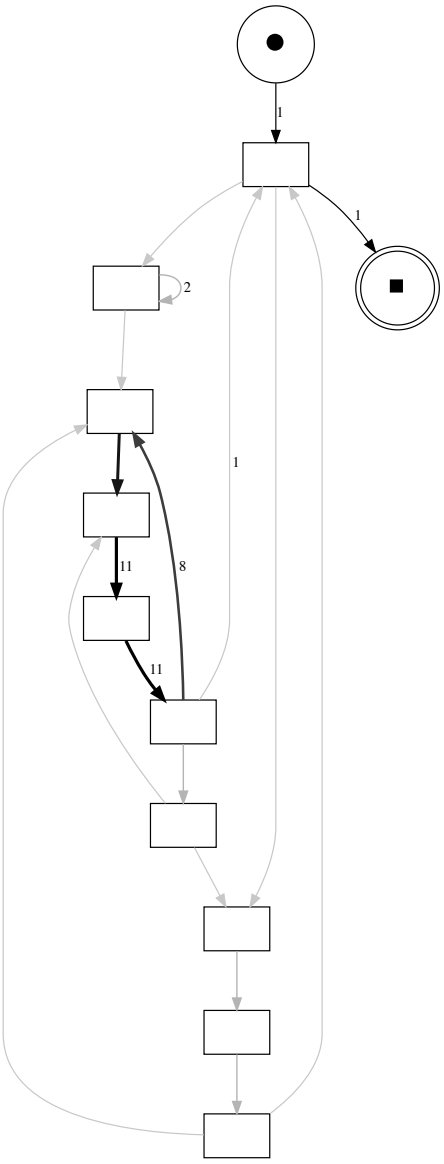

Figure 2: Entire song "As It Was" by Harry Styles as DFG of the part played by the left hand; activities represent measures, arc labels show frequencies.

we plan to realize different activity notions beyond measures, such as complete melodies, and thereby validate event abstraction techniques. Similarly, we experiment with our case notion to infer the process behavior of entire musical epochs possibly. With object-centric event logs, we could model the concurrent interplay of multiple instruments and the resolution of measures to notes. We plan to discuss our tool with music theorists to evaluate its applicability. It could give a new perspective for analyzing extensive music collections, e.g., the whole *Köchel catalog*, a collection of compositions by Wolfgang Amadeus Mozart. The current approach focuses on the modern Western notation of music, neglecting other notation systems of the world.

Similar formal notations exist for dance notation that could be similarly incorporated into our framework.

Process models as preprocessed scores could also serve as intermediary inputs to machine learning algorithms that output creative compositions or be used to explain them. Also, measure embeddings could be used instead of simple look-up tables, similar to word2vec. Given the large body of available music scores, our tool could help advance process mining techniques by providing countless (acoustically replayable and explainable) examples for training or be used in teaching to illustrate BPM concepts. Likewise, acoustically playable process models could facilitate the accessibility of process mining tools for visually impaired people.

Besides music scores, our approach is transferable to streaming settings by capturing live music. This would make the parallels to techniques such as conformance checking even more apparent, as slight variations like length and hitting a wrong note are recognizable as deviations.

References

[1] W. M. P. van der Aalst, J. Carmona (Eds.), Process Mining Handbook, Springer, Cham, 2022.

[2] A. Brysbaert, L. Bocchi, E. Tuosto, Relating archaeological chaîne opératoire and process mining in computer science, Archeologia e Calcolatori 23 (2012) 165–186. Publisher: Edizioni All'Insegna del Giglio.

[3] T. Li, M. Ogihara, G. Tzanetakis, Music Data Mining, CRC Press, Boca Raton, 2012. OCLC: 756680089.

[4] A. Baratè, G. Haus, L. A. Ludovico, Music Analysis and Modeling Through Petri Nets, in: R. Kronland-Martinet, T. Voinier, S. Ystad (Eds.), Computer Music Modeling and Retrieval, volume 3902, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 201–218. URL: http://link.springer.com/10.1007/11751069_19. doi:10.1007/11751069_19, series Title: Lecture Notes in Computer Science.

[5] M. Schuijer, Analyzing atonal music: pitch-class set theory and its contexts, Eastman studies in music, University of Rochester Press, Rochester, NY, 2008. OCLC: 213307961.

[6] M. S. Cuthbert, C. Ariza, music21: A Toolkit for Computer-Aided Musicology and Symbolic Music Data (2010). Publisher: International Society for Music Information Retrieval.

[7] A. Berti, S. J. van Zelst, W. van der Aalst, Process Mining for Python (PM4Py): Bridging the Gap Between Process- and Data Science, in: Proceedings of the ICPM Demo Track 2019 co-located with 1st International Conference on Process Mining (ICPM 2019), volume 2374, CEUR-WS, Aachen, Germany, 2019. URL: https://ceur-ws.org/Vol-2374/paper4.pdf.