



HAL
open science

A viewpoint approach to symbolic music transformation

Louis Bigo, Darrell Conklin

► **To cite this version:**

Louis Bigo, Darrell Conklin. A viewpoint approach to symbolic music transformation. 11th International Symposium on Computer Music Multidisciplinary Research (CMMR) 2015, Jun 2015, Plymouth, United Kingdom. hal-01798678

HAL Id: hal-01798678

<https://hal.science/hal-01798678v1>

Submitted on 23 May 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A viewpoint approach to symbolic music transformation

Louis Bigo¹

Darrell Conklin^{1,2}

¹ Department of Computer Science and Artificial Intelligence
University of the Basque Country UPV/EHU, San Sebastian, Spain

² IKERBASQUE, Basque Foundation for Science, Bilbao, Spain
{louis.bigo,darrell.conklin}@ehu.es

Abstract. This paper presents a general approach to the transformation of symbolic music. The method is based on viewpoints, which enable the representation of musical surfaces by sequences of abstract features. Along the transformation process, some of these sequences are conserved while some others are variable and can be replaced by generated ones. The initial piece is therefore seen as a template which is instantiated at each transformation. The method is illustrated in the paper with the particular case of transformations occurring at the harmonic level. New chord sequences are generated by sampling from a statistical model in a particular style. The pitch of the notes constituting the template piece are then transformed according to the generated chord sequence.

Keywords: harmonic transformation, viewpoints, computer-aided composition, harmonic analysis, music generation, statistical models, computational creativity

1 Introduction

Music generation methods can be divided into two broad categories [9, 11]. On one hand are *rule-based* methods that use hard coded rules and constraints for style emulation and algorithmic composition. On the other hand are *machine learning* approaches that generate musical objects by sampling from statistical models built from large corpora of music [6]. In this paper we propose a new approach to using statistical models for music generation, one guided by the transformation of a *template piece* from which intra-opus structural features are inherited. Generation by transformation has been investigated based on spatial representations [4] and audio content [2]. Some harmonic transformation methods have also been investigated to assist composition in the songwriting assistant system *Liquid Notes* [1]. A strong motivation of the transformational approach to music generation is to benefit from *conserved* high-level structures that are hard to generate. The generation can then be restrained to some *variable* musical objects, producing a transformation of the initial sequence that maintains its particular structural aspects.

An additional motivation for transformational approaches is to provide some tools to the composer along the creative process. A transformation system can

indeed be used at any step of the composition process to provide some *alternative* realisations of an abstract musical idea.

This work uses the symbolic music representation method of *viewpoints* [9, 8] for the formal representation of the conserved and variable aspects of a template piece. The general method is illustrated with the particular case of transformations occurring at the chord level. Pitches of notes are considered to be variable and are modified in order to fit with a generated chord sequence. Other aspects of the musical surface (rhythm, orchestration, etc.) are conserved and are left unchanged. The method presented here is used for the transformation of tonal sequences, and therefore requires a chordal analysis and key detection step.

The transformation algorithm can be summarized by three main steps. First, a harmonic analysis is performed on the template piece. Second, a new chord sequence is generated by random walk from a statistical model built from a corpus of chord sequences of a particular style. Finally, a new musical surface is produced by changing the pitch of every event of the template piece, constrained by the generated chords. Pitches are modified such that their harmonic function (e.g., chord note, passing note, etc.) and similarity with the template pitch sequence are conserved. The method therefore ensures that the register and the global melodic shape of the original piece are conserved.

This paper is structured as follows. The representation method of viewpoints is reviewed, with particular attention to harmonic viewpoints that are used in the steps of chord generation and pitch modification. Following this, the machine learning method for developing a statistical model of chord sequences is described, and the process for changing the notes of a piece is presented. Some transformation examples are illustrated using as a template an extract of Erik Satie’s *Gymnopédie No. 1* (1888).

2 A Viewpoint-based Method to Transform Musical Sequences

This section describes a general method for music transformation. A formalization of transformations based on viewpoint representations is introduced. The method is illustrated with the specific case of transformations occurring at the harmonic level which involves harmony-based viewpoints and generation of chord sequences in a particular style. Finally, a method to transform the template sequence according to a generated chord sequence is presented.

2.1 Viewpoint Representation for Transformations

The method of *viewpoints* is used to represent musical sequences. This representation method has already proven to be efficient in several fields like music prediction [9], music classification [8] and pattern discovery [7].

Musical sequences are represented at the surface level as sequences of events that have basic features including duration, onset time and additional values depending on the nature of the events (for example, note events include a pitch,

chord events include a chord symbol, etc.). A *viewpoint* is a function mapping events to more abstract derived features. The function is partial, therefore it may be undefined (\perp) for some events. An event e is abstracted by application of a viewpoint τ to produce the abstract feature $\tau(e)$.

The application of k viewpoints τ_1, \dots, τ_k to an event sequence e_1, \dots, e_n may be represented as a $k \times n$ solution array where location (i, j) holds the value $\tau_i(e_j)$. The upper table of Figure 1 illustrates such an array for a melody fragment extracted from the *Gymnopdie No. 1* of E. Satie. For better readability, only the notes of the melody part of the fragment (labeled by the events e_1, \dots, e_{10} on the upper staff) are represented with viewpoints. Values $\text{int}(e_i)$ and $\text{pc}(e_i)$ respectively correspond to the incoming pitch interval $\text{pitch}(e_i) - \text{pitch}(e_{i-1})$ and to the pitch class of the event $\text{pitch}(e_i) \bmod 12$. Though not represented by viewpoints, the accompaniment (lower staff) is kept on the figure to provide the reader the harmonic context, which is necessary to compute harmonic based viewpoints introduced below.

Transformations As previously mentioned, the notion of transformation requires a distinction between *conserved* and *variable* parts of an existing sequence. The process of transforming a sequence can then be seen as the task of modifying some of its describing viewpoint sequences while conserving some others. The choice of conserved and transformed viewpoints will be constrained by the dependencies between viewpoints. For example, modifying the pitch class of a note event will necessarily imply a modification of its pitch.

A strong advantage of this method, and also a major motivation of this work, is to enable the transformation of a musical sequence to be specified on higher musical levels (e.g., chords) than basic surface features (e.g., pitches). The process of transforming a musical sequence S can be described in 3 main steps (see Figure 2):

- represent S by a set of viewpoint sequences V ;
- produce an alternative set of viewpoint sequences V' by modifying some viewpoint sequences of V while conserving some others;
- generate a sequence of basic features S' that can be abstracted by the set of viewpoint sequences V' .

Note that because some viewpoint sequences in V are conserved, the transformed sequence S' has the same number of events as S . This general transformational approach is illustrated in the following with the specific case of transformation occurring at the harmonic level.

2.2 Harmony-based Viewpoints

This section introduces harmony-based viewpoints necessary to process harmonic transformations. In order to be computed, harmonic viewpoints require a preliminary harmonic analysis to be processed on the musical sequence.

key segmentation: D:maj

chord segmentation: G:M7, D:M7, G:M7, D:M7, G:M7

Notes: e_1 e_2 e_3 e_4 e_5 e_6 e_7 e_8 e_9 e_{10}

Note viewpoint	e_1	e_2	e_3	e_4	e_4	e_6	e_7	e_8	e_9	e_{10}
pitch	78	81	79	78	73	71	73	74	69	66
onset	1	2	3	4	5	6	7	8	9	12
duration	1	1	1	1	1	1	1	1	3	3
int	\perp	+3	-2	-1	-5	-2	+2	+1	-5	-3
pc	6	9	7	6	1	11	1	2	9	6
key	D	D	D	D	D	D	D	D	D	D
chord	G:M7	G:M7	D:M7	D:M7	D:M7	G:M7	G:M7	G:M7	D:M7	G:M7
hlab	c	ζ	ζ	c	c	c	ζ	c	c	c

Chord viewpoint	c_1	c_2	c_3	c_4	c_5
chord	G:M7	D:M7	G:M7	D:M7	G:M7
onset	0	3	6	9	12
duration	3	3	3	3	3
root	G	D	G	D	G
mode	maj	maj	maj	maj	maj
crm	\perp	P5	P4	P5	P4
cmm	\perp	maj→maj	maj→maj	maj→maj	maj→maj
crm \otimes cmm	\perp	$\langle P5, \text{maj} \rightarrow \text{maj} \rangle$	$\langle P4, \text{maj} \rightarrow \text{maj} \rangle$	$\langle P5, \text{maj} \rightarrow \text{maj} \rangle$	$\langle P4, \text{maj} \rightarrow \text{maj} \rangle$

Fig. 1. Five measures extracted from the *Gymnopédie* No. 1 of E. Satie. A harmonic segmentation of the excerpt is provided above the score. The two tables below the score provide a viewpoint representation of the sequence of 10 events constituting the melodic fragment, and its associated sequence of five chords provided by the harmonic analysis. The top part of each table shows the basic viewpoints, the bottom the derived viewpoints used in this paper.




Fig. 2. An illustration of the transformation process applied on a musical sequence S abstracted by a set of viewpoint sequences V . V' is a transformation of V . A transformed musical sequence S' consists of any sequence that can be abstracted by V' . The dotted line between S and S' illustrates the transformation that is made indirectly through V and V' .

Harmonic Segmentation Harmonic analysis includes as a first step the labelling of the sequence by chord and key segments. More specifically, a *chord segmentation* is a sequence of non-overlapping chord symbols, each labeled by a duration, that cover the time-line of a musical sequence. Additionally, a *key segmentation* is a sequence of keys, each labeled by a duration, covering the piece in the same way. The *harmonic segmentation* of a piece refers to the chord segmentation and the key segmentation resulting from the harmonic analysis of the piece. An harmonic segmentation is illustrated at the top of Figure 1. Though the harmony of the musical excerpt of Figure 1 is not ambiguous, in general there are no unique and exact methods for harmonic segmentation, in particular when inputs are MIDI files that do not include pitch spelling. These tasks are largely discussed within the music community and even when manually performed, they can produce different output depending on the analyst. Different methods trying to model this human cognitive ability have been investigated. These methods include an algorithm based on the spiral array [5], a dynamic programming approach [16] that processes chord and key segmentation based on Lerdahl’s tonal distance [13] and the Melisma system [17].

The transformation method presented in this paper requires a chord/key segmentation of the input sequence to compute harmony based viewpoints. This segmentation constitutes an additional input to the transformation. Whether it is manually performed or automatically computed by one of the previous systems does not impact the functioning of the transformation method. To generate the transformations discussed in Section 3, both the algorithm described in [16] and some manual harmonic segmentation were used.

Chord and Key Viewpoints A chord segmentation induces a viewpoint chord that returns for any note event e , the chord symbol of the chord segment in which e is included. The note viewpoint table of Figure 1 represents the chord viewpoint sequence (**chord**) associated with the melodic extract, which in that case includes values G:M7 and D:M7. Note that a more accurate harmonic analysis would typically depict the degrees of these chords as IV and I respectively, showing thus the Lydian quality of the sequence. Although it is not the case in the

harmonic transformations illustrated in this paper, the chord degrees could be conserved along the transformation by adding a chord degree viewpoint sequence to the conserved features of the transformation.

In this work, a note event is considered to be included within a segment if its onset is included in the segment. As a consequence, a note event that overlaps different segments will be systematically associated to the segment in which the event starts. Alternative segmentation strategies could be considered without affecting the functioning of the transformation method. In the same manner as chord segmentation, a key segmentation induces a viewpoint (**key**) that returns the key of the key segment that includes the event.

For an event e , the function chord_{pc} returns the pitch class set associated with the chord symbol $\text{chord}(e)$. Additionally, $\text{key}_{\text{pc}}(e)$ corresponds to the set of pitch classes that gathers all pitch classes composing the key $\text{key}(e)$. For example, in Figure 1, we have $\text{chord}_{\text{pc}}(e_6) = \{0, 4, 7, 10\}$ and $\text{key}_{\text{pc}}(e_3) = \{0, 2, 4, 5, 7, 9, 10\}$, respectively associated with the chord C:7 and the key F:maj.

Harmonic Label Viewpoint A contribution of the paper is the introduction of the viewpoint **hlab** that attributes a harmonic label to every note event of the template piece. For any event e , $\text{hlab}(e)$ is computed from the values $\text{pc}(e)$, $\text{chord}_{\text{pc}}(e)$ and $\text{key}_{\text{pc}}(e)$.

Though the notion of harmonic label can be defined in different ways, in particular depending on the musical style, a simple specification is proposed to illustrate the method. Three possible harmonic labels can be attributed to an event, depending on if its pitch belongs to its relating chord and key regarding the harmonic segmentation. More formally, we propose the set of harmonic labels $\{c, \kappa, o\}$ (c for “chord”, κ for “key” and o for “other”) with:

$$\text{hlab}(e_i) = \begin{cases} c & \text{if } \text{pc}(e_i) \in \text{chord}_{\text{pc}}(e_i) \\ \kappa & \text{if } \text{pc}(e_i) \notin \text{chord}_{\text{pc}}(e_i) \text{ and } \text{pc}(e_i) \in \text{key}_{\text{pc}}(e_i) \\ o & \text{if } \text{pc}(e_i) \notin \text{chord}_{\text{pc}}(e_i) \text{ and } \text{pc}(e_i) \notin \text{key}_{\text{pc}}(e_i) \end{cases}$$

On the example of Figure 1, we have $\text{hlab}(e_2) = \kappa$ because $9 \notin \{2, 6, 7, 11\}$ and $9 \in \{0, 2, 4, 6, 7, 9, 11\}$. Figure 3 provides two additional examples of melodic fragments and their harmonic label sequences. The first one is extracted from the Piano Concerto No. 21 of W.A. Mozart. The harmonic segmentation of this fragment is easily performed thanks to the accompaniment part, which is not represented on the figure. The second one is extracted from the jazz standard *Take the “A” train*. The harmonic segmentation of this fragment is taken from the original lead sheet. The above definition of **hlab** consists for every event in a mapping between the 12 pitch classes and the set of harmonic labels. A different specification that would require octave information of the events to specify their harmonic function would also be possible. The harmonic label attributed to each note depends on the output of the harmonic segmentation and on the variety of chord types and keys supported by the harmonic segmentation system.

The set of possible harmonic labels could include a larger variety of values than the three above, as for example the notion of fundamental within a chord.

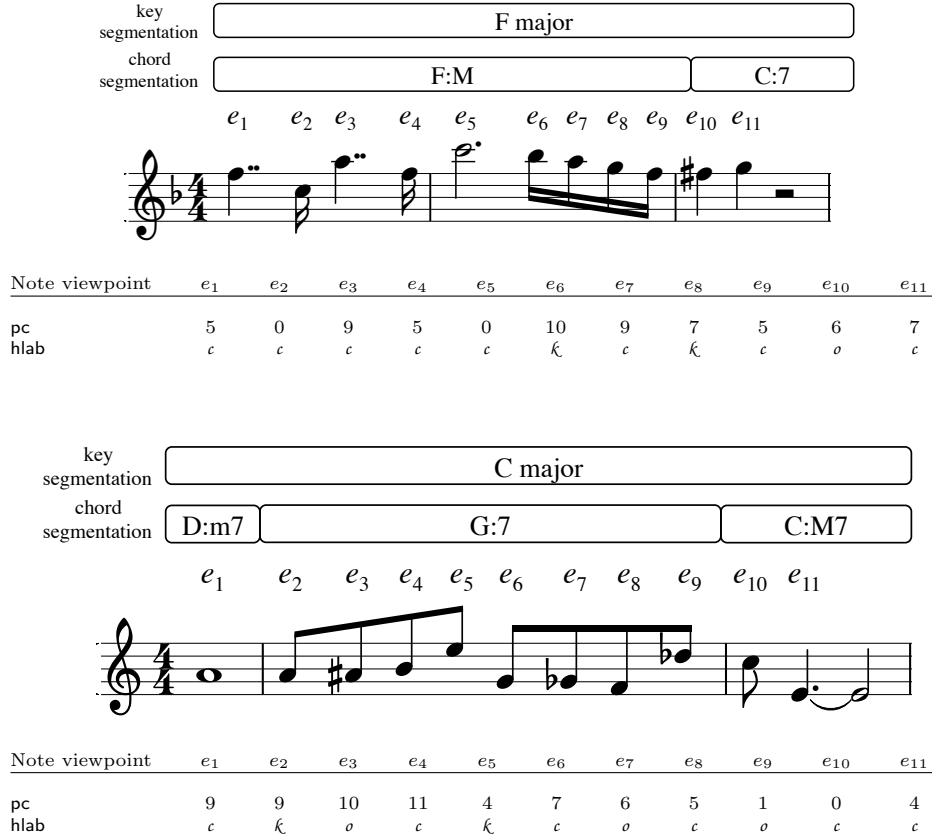


Fig. 3. Two examples of harmonic label sequences of melodic line fragments. The upper fragment is extracted from the second movement of the *Piano Concerto No. 21* of W.A. Mozart. The lower fragment is extracted from the jazz standard *Take the “A” train* from the pianist B. Strayhorn. For better readability accompaniment parts that are used to perform key and chord segmentation are not represented.

Some harmonic labels can be specific to some musical style, for example the notion of *blue note* in jazz. The definition of the set of harmonic labels impacts the precision of the harmonic description of the sequence. As illustrated in Section 3, this aspect acts as an interesting parameter in the transformation process.

2.3 Conserved Viewpoints

The harmonic transformations presented in this work consist in (1) generating a new chords sequence and (2) transforming a musical sequence regarding the newly generated chord sequence. Though the pitches of the original note events are transformed, their harmonic label, onset and durations are conserved. More

formally, an harmonic transformation of the musical sequence e_1, \dots, e_n is a sequence e'_1, \dots, e'_n that respects for every e' :

- $\text{onset}(e') = \text{onset}(e)$
- $\text{duration}(e') = \text{duration}(e)$
- $\text{hlab}(e') = \text{hlab}(e)$

While the above viewpoints are conserved, the others are transformed. Section 2.4 presents a method to generate a new chord sequence c'_1, \dots, c'_m that has the same length than the chord segmentation c_1, \dots, c_m and where each chord c'_i is attributed the same onset and duration than its corresponding original chord c_i . Section 2.5 presents a method to generate an harmonic transformation of a musical sequence regarding the new chord sequence c'_1, \dots, c'_m .

2.4 Chord Sequence Generation

The method for music generation using harmonic transformations relies centrally on the transformation of a template chord sequence into a new sequence. The task of chord sequence generation is stated simply as: given a statistical model over chord sequences, sample high probability sequences from the statistical model. This section describes the statistical modeling method used, and the corpus used to train the models.

A statistical model trained on viewpoint sequences from a corpus is used to generate new chord sequences. Here the method for using an abstract viewpoint to describe a first-order Markov model over chords is reviewed [8]. This method was also used recently to describe a statistical model for first-species counterpoint [12].

Let τ be a first-order viewpoint (i.e., computed from an event and its preceding event), and let $v = \tau(c_i | c_{i-1})$ be the feature assigned by τ to chord c_i , in the context of its preceding chord c_{i-1} . The probability $P(c_i | c_{i-1})$ of chord c_i following chord c_{i-1} can be written in the form

$$P(c_i | c_{i-1}) = P(c_i, v | c_{i-1}) = P(v) \times P(c_i | c_{i-1}, v)$$

with the first term $P(v)$ estimated as $c(v)/n$, where n is the number of chords in the corpus and $c(v)$ is the number of chords in the corpus having the feature v . To further reduce the number of parameters in the model simply to the possible values of τ , the second term $P(c_i | c_{i-1}, v)$ can be modelled with a uniform distribution over events having the feature v in the context of a given event c_{i-1} [8].

For a chord sequence c_1, \dots, c_m , the *cross-entropy* of the sequence according to the statistical model is the mean negative log probability of the sequence:

$$-\log_2 \prod_{i=2}^m P(c_i | c_{i-1})/m$$

To generate chord sequences an iterated random walk procedure is used [12]. The first chord is fixed to a chosen starting chord, then random walk is used to

generate a sequence of length m . This procedure is repeated multiple times with one of the low cross-entropy sequences retained. The random walk can be constrained to visit only tonal sequences that are composed by chords whose pitches belong to a unique tonality. Another possible constraint consists in generating sequences respecting a given structure (i.e., controlling how chords repeat along the sequence). The idea of constraining some chords along the generation process has also been explored by [14, 10]. The structure constraint is interesting in our context of transformation since it enables the conservation of the harmonic structure of the template piece. It will be illustrated in the second example of Section 3.

The statistical model of chords is created by compiling statistics from a chord corpus. The musical style of the corpus impacts the chord generation and should thus be chosen carefully depending on the transformations that aim to be processed. The *Academic* subsection of the 9GDB chord sequence corpus [15] has been chosen to perform the transformations discussed in Section 3. This corpus contains the three subsections of classical, baroque, and romantic, and includes 235 chord sequences, reaching to total of 13027 chords.

This corpus was transformed by two steps: first, all chord extensions and slashes (chord inversion specification) were removed and chords were truncated to major, minor, augmented, suspended and diminished triads; second, runs of the same triad were collapsed to just one occurrence.

2.5 Template Transformation

This subsection describes a method to transform the template sequence, given as a MIDI file, to fit with a generated chord sequence.

MIDI Transformation A MIDI file consists in a set of simultaneous tracks that can each provide a note event sequence e_1, \dots, e_n whose ordering corresponds to the ordering of their respective onset events within the MIDI track. However, two events e_i and e_{i+1} extracted from a polyphonic MIDI track can have the same onset time. As explained below, this property can have an important impact on a transformation.

For each track of the template MIDI file, a new sequence of events is computed with the method described below. Onsets and durations of the events are conserved but pitches are transformed. As a consequence, the MIDI file has the same structure before and after the transformation i.e., the same number of tracks, and the same number of events in each of these tracks.

Event Sequence Generation A notable property of the harmonic transformations presented in this work is to conserve the harmonic label of the notes. The transformation consists then in generating a note event sequence e'_1, \dots, e'_n such that $\text{hlab}(e'_i) = \text{hlab}(e_i)$ for every event e_i .

For any original event e_i , we call the *event candidate set* A_i the set of all possible events having the same onset and duration as e_i and having a pitch

respecting the harmonic label $\text{hlab}(e_i)$. For every event of the original sequence, an event candidate set is built. A transformed sequence then results from the choice of one event in each successive candidate set.

Event Sequence Selection Different strategies can be applied to select an event within a candidate set. To guide this selection, a score is calculated for any transformed sequence depending on its similarity with its original sequence regarding some arbitrary viewpoints.

Let $E' = e'_1, \dots, e'_n$ be the result of a transformation of a sequence $E = e_1, \dots, e_n$. The distance of the transformed sequence regarding a viewpoint τ corresponds to the mean distance, over the whole sequence, between $\tau(e_i)$ and $\tau(e'_i)$:

$$\delta_\tau(E', E) = \frac{\sum_{i=1}^n |\tau(e_i) - \tau(e'_i)|}{n}$$

For example, selecting an event sequence close from the original one in terms of pitches (i.e., minimizing δ_{pitch}) will tend to conserve the register of the original pitch sequence. On the other hand, selecting a sequence similar in terms of pitch interval (i.e., δ_{int} has a low value) will maintain the global pitch shape of the sequence. In the case of a polyphonic MIDI track, minimizing δ_{int} will tend to approximate both horizontal and vertical intervals of the original track.

This method only holds for numerical viewpoints. For non-numerical viewpoints, a specific notion of distance between $\tau(e_i)$ and $\tau(e'_i)$ would have to be defined. Note that τ should not belong to the viewpoints that are conserved along the note transformation process, otherwise, $\delta_\tau(E', E)$ would be zero. In the case of harmonic transformations presented here, it can be neither chord, hlab, onset nor duration.

In the following, we propose different strategies to generate sequences giving a low value to $\delta_\tau(E', E)$:

Greedy Algorithm This strategy consists in selecting for each e_i a pitch in A_i that minimizes the local distance $|\tau(e_i) - \tau(e'_i)|$. Though very efficient, this algorithm provides only one solution that is not guaranteed to be optimal.

Viterbi Algorithm This strategy consists in finding the sequence that minimizes $\delta_\tau(E', E)$. This task is achieved with dynamic programming [3]. If $\tau(e_i)$ is a 0-order viewpoint (i.e., $\tau(e_i)$ only requires e_i to be computed), the Viterbi algorithm will return the same sequence as the greedy algorithm. This is the case for the viewpoint $\text{pitch}(e_i)$. However, first-order viewpoints (e.g., $\text{int}(e_i)$ takes into account e_{i-1}) will benefit from the Viterbi algorithm to compute the sequence that minimizes $\delta_\tau(E', E)$.

Random Walk A score is attributed to every event e of a candidate set A_i . This score is inversely proportional to the distance $|\tau(e_i) - \tau(e'_i)|$. An event is then randomly sampled from the set of candidate events, according to their relative

score. This strategy has the advantage of providing a large number of solutions in efficient time.

3 Analysis of Transformed Sequences

This section illustrates our method by presenting and discussing different harmonic transformations of a template sequence corresponding to the five measures extracted from Erik Satie’s *Gymnopédie No. 1* that are illustrated in Figure 1. The chord sequences are generated from a statistical model built from the *Academic* subsection of the 9GDB chord sequence corpus [15] with chords truncated to major, minor, augmented, suspended and diminished triads as explained in Section 2.4. As a consequence, generated sequences only include these types of chords in these transformation examples. Generating more complex chords (e.g., including sevenths) would be possible but would require the statistical model to handle a more sophisticated chord representation.

Obviously these examples do not aim at producing anything comparable to the original sequence from an aesthetic point of view. In particular, the harmonic singularity of this piece, sometimes considered as post-tonal, did not influence these transformations since only onsets and durations of the original chords are conserved along the process. Furthermore, we choose to exemplify an harmonic transformation on an extract being both well-known and rhythmically simple, in order to illustrate more intuitively the effects of the transformations¹.

3.1 The Template Sequence

Figure 1 displayed at the beginning of the paper illustrates the score of the template sequence and the viewpoint representation of its melodic part. The events of the accompaniment part belong to a separate track which is not represented on the table for better readability. However, this track is transformed with the same method.

The preliminary harmonic segmentation of this template piece is not ambiguous. As shown under the score, it consists in an alternating of the chords G:M7 and D:M7 in the key of D major. The harmonic label sequence only provides values c and ξ which means that the fragment does not include any note outside the key of D major.

A less accurate harmonic analysis might have returned G:maj and D:maj for the two chords. As a consequence, the pitch class $F\sharp$ of e_1 would not be considered of being part of the chord $\text{chord}(e_1)$ and $\text{hlab}(e_1)$ would then have the value ξ instead of c . A transformation of the sequence would then assign to e'_1 a pitch not included in the generated chord $\text{chord}(e'_1)$. This example typically illustrates how the harmonic analysis method impacts the transformation process. Interestingly, it does not seem obvious that the quality of the transformation is proportional

¹ The original excerpt and transformations are available at the address : <https://soundcloud.com/harmonictransformations>.

Note viewpoint	e'_1	e'_2	e'_3	e'_4	e'_4	e'_6	e'_7	e'_8	e'_9	e'_{10}
pitch	78	83	78	79	71	73	74	73	69	64
pc	6	11	6	7	11	1	2	1	9	4
key	D	D	D	D	D	D	D	D	D	D
chord	D:maj	D:maj	E:min	E:min	E:min	A:maj	A:maj	A:maj	D:maj	A:maj
hlab	c	κ	κ	c	c	c	κ	c	c	c

Chord viewpoint	c'_1	c'_2	c'_3	c'_4	c'_5
chord		D:maj	E:min	A:maj	D:maj
				D:maj	A:maj

Fig. 4. A transformation of the extract of Figure 1.

to the accuracy of the harmonic analysis. In this example, forcing e'_1 to be out of the chord would produce a larger variety of transformed sequences, which can be interesting from a creativity point of view.

3.2 Transformation A

Figure 4 illustrates a first harmonic transformation of the fragment. The chord generation process has been performed with a filtration of tonal sequences, as explained in Section 2.4. The generation returns a low entropy triadic sequence in the key of D:maj (see bottom line of Figure 4).

As explained in Section 2.3, the viewpoint sequence `hlab` is conserved through the transformation. The pitches of the events e'_1, \dots, e'_{10} are generated by computing the greedy solution minimizing local distance $|\text{pitch}(e_i) - \text{pitch}(e'_i)|$ tending thus to approximate the value of the pitches of the template sequence. The same algorithm has been applied to generate the event sequence constituting the accompaniment part. A notable property of this transformation is that some chords in the accompaniment part of the resulting score include fewer notes than in the template. This is the case for the chords appearing in bars 2, 3 and 5. This is due to a side effect of the strategy consisting in approximating the pitch value of the template notes: some simultaneous notes of the template piece have their pitch transformed into the same new pitch, producing identical events. This effect can be handled by approximating intervals between events rather than pitches, as proposed in the next example.

Note viewpoint	e'_1	e'_2	e'_3	e'_4	e'_4	e'_6	e'_7	e'_8	e'_9	e'_{10}
pitch	79	81	79	77	72	72	74	76	69	67
int	\perp	+2	-2	-2	-5	0	+2	+2	-7	-2
pc	7	9	7	5	0	0	2	4	9	7
key	C	C	C	C	C	C	C	C	C	C
chord	C:maj	C:maj	F:maj	F:maj	F:maj	C:maj	C:maj	C:maj	F:maj	C:maj
hlab	c	κ	κ	c	c	c	κ	c	c	c

Chord viewpoint	c'_1	c'_2	c'_3	c'_4	c'_5	
chord		C:maj	F:maj	C:maj	F:maj	C:maj

Fig. 5. A transformation of the extract of Figure 1.

Another observation on the accompaniment part of this transformation is that the root note of the first chord D:maj is an F \sharp which puts this chord in an inverted position contrary to the template in which all chords are in root position. Maintaining chord positions along transformation can be handled by adding the notion of chord fundamental in the set of available harmonic labels as illustrated in the next transformation example.

3.3 Transformation B

Figure 5 illustrates a second harmonic transformation of the fragment. The structure of the template chord sequence is conserved by applying a filter along the chord generation process as mentioned in Section 2.4. As a consequence, the generated sequence consists, as the original one, in two alternating chords. The pitches of the events e'_1, \dots, e'_n constitute the optimal solution (computed with the Viterbi algorithm mentioned in Section 2.5) minimizing the distance δ_{int} approximating thus the global pitch shape of the original sequence. Unlike in the previous transformation, simultaneous events will unlikely be transformed into identical events (i.e., having the same pitch) with this strategy. As a result, the number of notes in the chords in the accompaniment part is conserved. Furthermore, this transformation has been made while taking into account an additional harmonic label in the set presented in Section 2.2. This harmonic label specifies whether the pitch class of an event corresponds to the fundamental of its associated chord. This modification does not impact the transformation

of the melodic part of the template because this part does not include any note whose pitch corresponds to the fundamental of its associated chord. However, the accompaniment part is affected and every chord is voiced in root position.

4 Conclusions

This paper presented a viewpoint approach to transform symbolic musical sequences. The method has been illustrated with the particular case of transformations occurring at the harmonic level and two examples of harmonic transformation have been discussed. Harmonic transformations can be controlled by a large set of variable components including the harmonic analysis system and the set of harmonic labels, the style and the complexity of the chords constituting the corpus which is learned by the statistical model, and the algorithm used to transform note event sequences.

The generic aspect of the viewpoint approach suggests a wider range of musical transformations that constitute future perspectives of this research. It is planned to explore rhythmic transformations in which onset and duration viewpoints sequences would be modified. The possibility to add or remove events from the original sequence is also part of future work.

Acknowledgments The authors thank Dorien Herremans for valuable discussions and collaboration on this research. This research is supported by the project Lrn2Cre8 which is funded by the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 610859.

References

1. Liquid Notes <http://www.liquid-notes.com>.
2. X. Amatriain, J. Bonada, A. Loscos, J. L. Arcos, and V. Verfaillie. Content-based transformations. *Journal of New Music Research*, 32(1):95–114, 2003.
3. R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, USA, 1 edition, 1957.
4. L. Bigo, D. Ghisi, A. Spicher, and M. Andreatta. Spatial transformations in simplicial chord spaces. In *Proceedings of the joint International Computer Music Conference — Sound and Music Computing*, pages 1112–1119, Athens, 2014.
5. E. Chew. The spiral array: An algorithm for determining key boundaries. In *Music and artificial intelligence*, pages 18–31. Springer, Lecture Notes in Computer Science Volume 2445, 2002.
6. D. Conklin. Music generation from statistical models. In *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in the Arts and Sciences*, pages 30–35, Aberystwyth, Wales, 2003.
7. D. Conklin. Discovery of distinctive patterns in music. *Intelligent Data Analysis*, 14(5):547–554, 2010.
8. D. Conklin. Multiple viewpoint systems for music classification. *Journal of New Music Research*, 42(1):19–26, 2013.

9. D. Conklin and I. Witten. Multiple viewpoint systems for music prediction. *Journal of New Music Research*, 24(1):51–73, 1995.
10. Arne Eigenfeldt and Philippe Pasquier. Realtime generation of harmonic progressions using controlled markov selection. In *Proceedings of ICCM-X-Computational Creativity Conference*, pages 16–25, 2010.
11. J. D. Fernandez and F. J. Vico. AI Methods in Algorithmic Composition: A Comprehensive Survey. *Journal of Artificial Intelligence Research*, 48:513–582, 2013.
12. D. Herremans, K. Sørensen, and D. Conklin. Sampling the extrema from statistical models of music with variable neighbourhood search. In *Proceedings of the joint International Computer Music Conference — Sound and Music Computing*, pages 1096–1103, Athens, 2014.
13. F. Lerdahl. *Tonal pitch space*. Oxford University Press, 2001.
14. F. Pachet and P. Roy. Markov constraints: steerable generation of Markov sequences. *Constraints*, 16(2):148–172, 2011.
15. C. Pérez-Sancho, D. Rizo, and J-M. Iñesta. Genre classification using chords and stochastic language models. *Connection Science*, 20(2&3):145–159, 2009.
16. T. Rocher, M. Robine, P. Hanna, and R. Strandh. Dynamic chord analysis for symbolic music. In *Proceedings of the International Computer Music Conference*, Montreal, Quebec, Canada, 2009.
17. D. Sleator and D. Temperley. The Melisma music analyzer. *Available online at www.link.cs.cmu.edu/music-analysis*, 2001.