Non-Conformant Harmonization: the Real Book in the Style of Take 6

François Pachet, Pierre Roy
Sony CSL
Paris, France
pachetcsl@gmail.com

Abstract

We address the problem of automatically harmonizing a leadsheet in the style of any arranger. We model the arranging style as a Markov model estimated from a corpus of non-annotated MIDI files. We consider a vertical approach to harmonization, in which chords are all taken from the arranger corpus. We show that standard Markov models, using various vertical viewpoints are not adapted for such a task, because the problem is basically over constrained. We propose the concept of *fioriture* to better capture the subtleties of an arranging style. Fioritures are ornaments of the given melody during which the arranging style can be expressed more freely than for melody notes. Fioritures are defined as random walks with unary constraints and can be implemented with the technique of Markov constraints. We claim that fioritures lead to musically more interesting harmonizations than previous approaches and discuss why. We focus on the style of Take 6, arguably the most sophisticated arranging style in the jazz genre, and we demonstrate the validity of our approach by harmonizing a large corpus of standard leadsheets.

Introduction

Automatic harmonization has been addressed for decades by computer music research (see Steels, 1986 for an early attempt at machine-learning of harmonization and Fernandez and Vico, 2013 for a survey). One reason for the success of this problem in the research community is that it can be considered, in first approximation, as a *well-defined problem*, a crown jewel in computer music. Automatic harmonization denotes in practice many different problems, depending on the nature of the input (melody, chord labels, bass, song structure given or not) and of the output (chord labels, chord realizations, contrapuntal voices), the constraints concerning the nature of the targeted harmonization (number of voices) and the way the targeted style is modeled (programmed explicitly or learned from examples). A widely studied variant of the automatic harmonization problem is the generation of a four-part (or more) harmonization of a given melody. Such a problem has been tackled in a variety of contexts, though mostly for classical music, Bach chorales in particular, and using virtually all the technologies available including rules, functions (Koops et al. 2013), grammars, constraints (Anders and Miranda, 2011), and statistical models of all types (Paiement et al. 2006).

Today, there are many approaches that work satisfactorily to produce harmonizations in the Classical style with reasonable musical quality. It is remarkable that automatic harmonization has achieved such a status of well-definedness that many papers in this domain consist in variations of existing algorithms, with little or no musical output (a sign, probably of the maturity of the field). However, there is no system, to our knowledge, that is able to produce truly musically interesting harmonizations, at least for the ears of musically trained listeners such as the first author of this paper. In the context of computational creativity, we claim that there are two problems with the current state of the art which limit their quality, and therefore their possibility for generating creative outputs: *excess of conformance* and *excess of agnosticism*.

Conformance. Automatic harmonization has so far been envisaged solely under the viewpoint of harmonic conformance: the main criterion of success is that the generated material has to conform to the harmonic constraints of the problem. For instance, a harmonic label of C minor (either imposed or inferred from, say, a soprano) should produce chord realizations that *conform* to C minor, for instance, chords composed of important notes of the scale. Conformance yields indeed a well-defined measure to evaluate systems, because there are well-defined harmonic distances (see Section Harmonic Distance), but tends to go in the way of creativity, since the best a system can do is to paraphrase harmonic labels. Such a skill can be impressive for non-musicians, but not for experts. Consequently, many harmonization systems give the impression that they are essentially filling the blanks (inner voices) with correct but uninteresting musical excipients. This is sometimes referred to as the “correct” versus “good” problem, but in fact, such harmonizers are basically unable to produce interesting solutions, because of excess in conformance.

Agnosticism (excess of generality). Most works, with the exception of (Ebcioglu, 1986), attempt to model a given style using general methods (such as Markov models, rules, etc.). General methods can be good in general, but are rarely very good in particular. Similarly to the famous “glass ceiling” problem occurring in MIR (Casey et al., 2008), there seems to be a glass ceiling concerning the musical quality of automatic harmonization. In our view this is caused by the use of too general methods and by the absence of consideration for the details of what makes a specific style interesting or creative. Most often, these details are not captured by general methods.

In this study, we focus on the harmonization style of the American six-voice *a cappella* band Take 6. Take 6 is the most awarded vocal group in history. Since their first two albums (Take 6, 1988; 1990) they renewed the genre of gospel barbershop-like harmonization by pushing it to its harmonic and vocal limits. Their style of arranging is considered unanimously as extraordinarily inventive, recognizable, and very difficult to imitate. Even the transcription of their performances is a very difficult task that only harmony experts can perform correctly (see Section Acknowledgements). Most of their works consist in 6-voice note-to-note harmonization of traditional songs, with many dissonances and bold voice movements typical of jazz big bands. The creativity of Take 6, if any, consists precisely in the use of those dissonances and digressions. Of course, their style and specificity is arguably also dependent on the quality of the singing voices (notably the bass), but this dimension is outside the scope of this paper, and we consider here only the symbolic aspects of their arranging style.

Most knowledgeable listeners of Take 6 enjoy “wow” effects due to their spectacular use of harmonic surprises. Figure 1 shows an excerpt of a harmonization by Take 6 of the traditional “Hark the Herald Angels Sing”. Figure 2 shows an estimation of the corresponding excerpt of the leadsheet (end of section A). It can be seen clearly that the chords used to harmonize the note B♭ do not conform to the expected harmony of B♭ major: although the performance of Take 6 are not labeled, we can estimate the last realization of the B♭ as an instance of a *C7dim9#11* (*C E G Bb Db F#*), which is very far from the expected Bb major scale, or of any scale close by (such as relative minors). Such a harmonic surprise is typical of the style of Take 6. By definition, conformant methods in automatic harmonization are not able to capture this kind of knowledge, especially from non-labeled training data.

Our goal is to produce six-voice harmonization in that style that triggers the same kinds of “wow” effects as the originals. The key idea of our approach is that most wow effects are obtained by *non-conformant harmonizations*, i.e., harmonizations that do not conform to the harmonic labels of the original leadsheet, but stay within well-defined constraints. The technical claim of this paper it that the technology of Markov constraints (Pachet et al., 2011) is particularly well suited for such a task, thanks to the possibility of generating creative sequences within well-defined constraints.

Problem Statement

The problem we address constitutes a variation on standard harmonization problems such as melody or bass given. It can be defined in terms of inputs/outputs as follows:

Inputs:

- A leadsheet representing the target melody to harmonize, as well as chord labels in a known syntax (i.e., we know their pitch constituents),
- A harmonization style represented by a set of non-annotated scores containing polyphonic content. No annotation of these scores is needed. In practice, arbitrary MIDI files may be used, including files without a fixed tempo coming from, e.g., recordings of real-time performances.

The expected output is a fully harmonized score, in the given style, i.e. a polyphonic score that maintains the the soprano of the leadsheet, and whose harmonies fit with the leadsheet chord labels.

Figure 1. Example of a typical non-conformant harmonization by Take 6. Harmonies used (estimated from the score) go from Bb (which conforms to the leadsheet) to a surprising, non-conformant C7dim9#11 (Transcription by A. Dessein).

Figure 2. Extract of a leadsheet for “Hark the Herald Angels Sing” (end of section A). The last B*b* is supposed to be harmonized in Bb (shortcut for Bb major).

Musically, the goal is to produce a harmonization that is *reminiscent* of the style, i.e., such that knowledgeable listeners can recognize the authors. However this is not a well-defined problem, for several reasons: listeners may not recognize a style because they do not know the arranger well enough, or because they give more importance to the sound than to the notes, or for many other reasons, including that the arranger may not have any definite style *per se*. In this paper, we do not attempt to solve the harmonization problem in many styles (though the system can, as exemplified in Section Applications to Other Styles). Rather, we attempt to convince ourselves, as knowledgeable Take 6 listeners, that our system grasps some of their subtle arranging tricks and reproduce them in unknown situations. A scientific evaluation of the system based on style recognition is in progress but is not the subject matter of this paper.

Corpora Used

The experiments we describe use a comprehensive database of jazz leadsheets described in (Pachet et al., 2013). For each leadsheet we have a melody (monophonic sequence of notes) and chord labels. For each chord label, the database provides the set of pitch-classes of the chord, in ascending order (that is, the formal definition of the chord, not its realization). In this study we used the Real Book (illegal edition), the most widely used jazz fake book. The Real book contains about 400 songs, 397 of which are parsed correctly (a few songs with no harmony or no melody are ruled out for instance).

For the harmonization style, we have selected a number of composers including classical ones (Wagner, Debussy, etc.) and jazz (Take 6 notably, and Bill Evans). Each composer is represented by a set of MIDI files of some of their compositions or performances. All MIDI files have been found on the web, except the 10 MIDI files of Take 6 that were provided to us by a human transcriber (A. Dessein). The Take 6 MIDI files are of excellent quality (i.e., there are virtually no transcription errors). The other MIDI files are of varying quality. Some of them correspond to actual scores (Wagner), others to performances (Bill Evans). IN order to cope with the diversity of tonalities and pitch ranges encountered in the leadsheet melodies, we have transposed systematically the corpus in all 12 keys.

Homophonic Harmonization

The approach we follow consists in considering the harmonization problem as a *vertical* problem, as opposed to voice-leading approaches (such as Whorley et al., 2013), and following an older tradition initiated by (Pachet and Roy, 1995) on constraint-based 4-voice harmonization in the Classical style. To compensate for the monotony of strict vertical harmonization, we complement this step by a smoothing procedure that somehow reestablishes voice-leading *a posteriori* from the vertical skeleton structure, by joining contiguous notes with the same pitch. This second step is completely deterministic, and the central issue we address is the production of the chordal skeleton.

Before describing the harmonization process, we introduce a measure of *harmonic conformance*, which is at the core of the whole process.

Harmonic Conformance

Because the scores of arrangers are not labeled, we need a way to relate chord realizations found in the arranger corpus to chord labels of a leadsheet. In order to avoid the pitfalls of chord recognition (which works well for simple chords, but much less for the complex chords as found in jazz), we use a simple but robust measure of the harmonic conformance between unlabeled chords. This measure, called ε-conformance is based on pitch class histograms.

For any chord realization C_i, i.e., a set of MIDI pitches, we build a pitch class histogram as an array of 12 integers, where each integer represents the number of occurrences of the corresponding pitch class in the chord (starting with C up to B), normalized by the total number of pitches. For instance, the circled chord in Figure 1 has a pitch-class frequency count $f = [0,1,1,0,1,0,1,0,0,1,0,0]$. The histogram is the frequency count divided by its module $h_i = \frac{f_i}{\sum_{k=1}^{12} {f_k}^2}$. The harmonic distance between two chords C_1 and C_2 can then be defined as the scalar product of the pitch class histograms:

$$D(C_1, C_2) = 1 - \sum_{k=1}^{12} h_1^k \times h_2^k.$$

where h_1(resp. h_2) is the pitch-class histogram of chord C_1 (resp. C_2). Such a distance takes its values in $[0, 1]$.

In practice, this distance enables us to categorize chord realizations appearing in the arranger corpus with regards to a given chord label. For each chord label, we can define an ideal prototype consisting of its pitch class definition, and then consider the ball centered on this ideal prototype of radius ε. ε represents the “harmonic conformance” of a chord realization to a chord label. Increasing values of ε provide increasingly large sets of chords, that are more or less conformant to the label. Figure 3 shows examples of chords at various distances to “C 7” for various values of ε in the Take 6 corpus.

Another way to relate chord realizations to chord labels is to consider the *best match* for a given corpus: the chord in the arranger corpus with the minimal harmonic distance to the ideal realization of the label. We then consider the ball centered around this best match, of radius ε. In any case, pitch class histograms provide us with a robust way to fetch chord realizations for any chord label, in non-annotated corpora.

Unary Markov Constraints

Equipped with a harmonic distance, we can generate new chordal skeletons. The idea is to estimate a Markov model of the sequences of chord realizations from the arranger corpus. The leadsheet (soprano movement and chord labels) is represented as a set of unary constraints holding on the sequence to generate. The framework of Markov constraints (Pachet et al., 2011), is precisely designed to handle such cases, and provides an efficient algorithm to generate those sequences, as well as a guarantee that all sequences satisfying the constraints will be found, with their correct probability in the original model. Solving a Markov constraint problem is strictly equivalent to sampling the sequences in the space of solutions. Each sequence $s = s_1, .., s_n$, has a probability $p(s) = p(s_1) \times \prod_{i=1}^{n-1} p(s_{i+1}|s_i)$ according to the considered Markov model (see next section). The unary Markov constraint algorithms guarantee that all sequences satisfying the constraints are drawn with their probability in the original model.

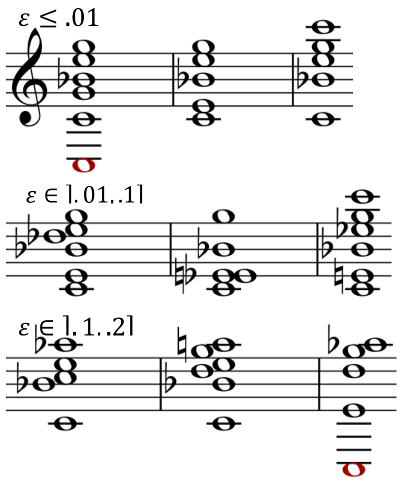

Figure 3. Various chord realizations from the Take 6 corpus for several values of ε (0.01, 0.1 and 0.2), representing increasing harmonic distance to a *C 7* chord label. As ε increases, more notes outside of the legal notes of C 7 (C, E, G, Bb) are added. For $\varepsilon = 1$ (maximum distance) all possible chords of the corpus are considered. In practice, reasonable, conformant realizations lie within a distance of about .15.

Viewpoints

Such a process raises an important issue concerning the choice of the viewpoint, i.e. the actual data used to estimate the Markov model. The most demanding viewpoint is the actual set of notes (Midi pitches) of the chord. This is called here the *Identity* viewpoint, since it contains all the information we have on a chord. Degraded viewpoints are also considered: *BassTenorSoprano* is the viewpoint consisting of the bass, tenor and soprano pitches (and ignoring the others). We define similarly the *BassSoprano* and *Soprano* viewpoints. For the sake of comparison, we also introduce the *Constant* viewpoint, which assigns a constant value to any chord (and serves as a base line for our experiments). Note that we do not consider duration information, as we do not want to rely on the quality of the MIDI Files.

Of course there is a tradeoff here between 1) harmonic conformance, represented here by ε, and 2) style conformance, which manifests itself by the presence of chord transitions that actually occurred in the corpus. Such a tradeoff between adaptation and continuity is not novel, and has been studied in automatic accompaniment (Cabral et al., 2006; Marchini and Purwins, 2010). In our context, it is formulated as a tradeoff between ε and viewpoint selectiveness. The most demanding viewpoint generate chord sequences that sound more natural in the given style, since they replicate actual transitions of chord realizations occurring in the corpus. However, such chord transitions will generate a sparse Markov model. The consequence is that only a very small number of leadsheets can be harmonized in that way for small values of ε. By degrading the viewpoints, more transitions will be available, so smaller (more conformant) values of ε can be considered.

Harmonizing the Real Book

In order to illustrate the harmonic conformance / viewpoint tradeoff, we describe a basic experiment that has, to our knowledge, never been conducted, at least on such a scale. For several values of ε we study the sparsity of the four viewpoints introduced above, by counting how many songs from the Real Book can be harmonized entirely with the viewpoint.

More precisely, for each leadsheet taken from the Real Book (397), we build a Markov Constraint problem consisting of the following constraints:

- Generate a sequence of chord realizations taken exclusively from the Take 6 corpus, transposed in all 12 pitches (variable domains),
- Each note of the leadsheet is harmonized by one chord realization (homophonic note-to-note harmonization),
- Transitions between 2 chord realizations c_i and $c_{i+1}$are all Markovian for the considered viewpoint, i.e. $p(c_{i+1}|c_i) > 0$,
- Each chord c_{i+1} has a soprano which is the leadsheet note
- Each chord realization c_{i+1} must be ε-conformant to the corresponding leadsheet chord label, for the chosen value of ε.

These constraints can all be implemented as a unary Markov constraint problem. The experiment consists in counting, for each value of ε in $[0, 1]$ and for each of the four viewpoints how many songs from the Real Book can be fully harmonized. The results are presented in Figure 4. It can be seen clearly that with non-trivial viewpoints (i.e. all viewpoints but soprano), solutions are found only for high values of ≤ε. For those values, harmonic conformance is lost. Only the basic *Soprano* viewpoint leads to many solutions (160, a value insensitive to ε). It can be noted that the *Constant* viewpoint (a trivial viewpoint that consists in basically removing the Markovian constraint), solutions are found for 262 songs. This means that there are 102 songs for which the *Soprano* viewpoint does not lead to any solution, for any value of ε. This corresponds to songs that contain pitch transitions that never occur between two consecutive realized chords in the Take 6.

It is important to note here that when no solution is found for a given leadsheet / viewpoint combination / value of ε, this does *not* necessarily implies that the leadsheet contains a transition for which there is no match in the corpus (for the given viewpoint). It means that there is no *complete* solution, i.e. transitions compatible with each other so as to make up a complete solution sequence.

This experiment shows clearly that harmonic conformance is somewhat incompatible with precise Markov models of chord realizations, for a realistic corpus (Take 6) on a realistic test database (the Real Book). However, we can use the Soprano viewpoint as a basis for producing interesting harmonization of most "reasonable" leadsheets, with a clear control on harmonic conformance.

Figure 5, Figure 6 and Figure 7 show homophonic harmonization of the four first bars of Giant Steps with various values of ε. It can be noted that while harmonic conformance can be used as a parameter to generate more or less conformant realization, the results are academically correct, but rarely very interesting musically. The style of the arranger is hard to recognize, because there are not enough actual transitions that are being reused from the corpus. The control of harmonic conformance can generate surprises, but at the price of losing the essence of the style.

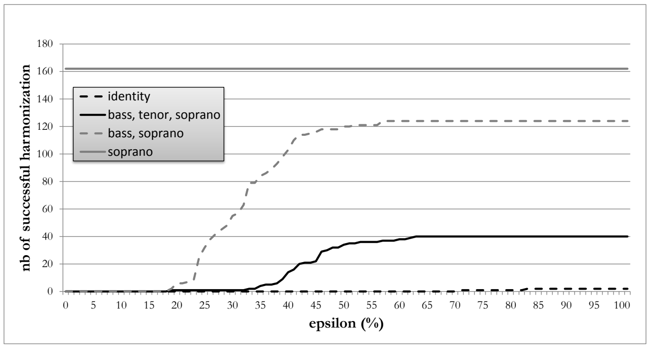

Figure 4. Graph showing the number of successful harmonization from the Real Book (illegal edition) using a Markov model of chord realizations, and various viewpoints of decreasing precision (identity, bass/tenor/soprano, bass/soprano, soprano).

Figure 5. The beginning of Giant Steps harmonized with a value of $\varepsilon \in [0, 0.01]$. All realizations come from the Take 6 corpus satisfy exactly the chord labels. The overall harmonization is conformant but not very interesting.

Figure 6. The beginning of Giant Steps with $\varepsilon \in [0, 0.2]$. The chords are less conformant and more interesting, but the whole harmonization still lacks surprise.

Figure 7. The beginning of Giant Steps with $\varepsilon \in [.3, .4]$. Chords are clearly farther away from the label, while retaining some flavor of the labels. However the decrease in harmonic conformance is musically not very interesting.

In order to express the harmonization style more clearly, and simultaneously bring creativity in the harmonization process, we introduce the concept of Fioriture.

Fioritures as a stylistic device

The idea of fioriture comes from a simple observation of polyphonic scores written by masters: It is difficult to be inventive on short duration notes. However, long notes raise opportunities to express a style: the longer a note is, the more possibilities of invention the arranger has. In the context of leadsheet based harmonization, we therefore introduce the concept of fioriture as a free variation, in the style of the arranger, occurring exactly during a *long* note, and making sense with its context.

A Simple Fioriture Example

We illustrate the concept of fioriture on a simple example. The task is to harmonize the melody shown in Figure 8: two notes with simple chords labels (both notes belong to the chord triads).

Figure 8. A simple melody to harmonize with fioritures.

This melody can be harmonized homophonically as described above, as illustrated in Figure 9.

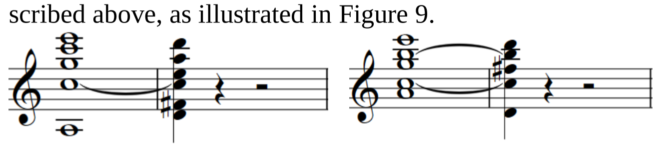

Figure 9. Two homophonic harmonization of the melody in Figure 10, with $\varepsilon \in [0, .1]$and $\varepsilon \in [.1, .2]$ respectively. A higher value of ε, the second one is more jazzy with a 9th added to the first chord and a 6th to the second one.

We can generate here a fioriture on the first note, since its duration is 4 beats. The Markov constraint problem corresponding to this fioriture is the following:

- First, select a rhythm for a note starting on the first beat of a 4/4 bar, and lasting 4 bars (rhythm selection is described in the next section). Let n be the number of notes, we generate $n + 1$ chord realizations to include the chord on the following note (here a D).
- The domain of the first chord contains only chords whose soprano is the first melody note (here, A).
- We can choose here a demanding viewpoint such as the *identity* viewpoint because in most cases the constraints above are not too hard.

Figure 10 shows various solutions, with increasing number of notes in the fioriture. It should be noted that all fioritures start from a soprano *A* on a *Amin* chord and end on a soprano *D* on a *D7* chord. However, some of them, in particular the last ones, deviate substantially from the chord labels. In short, they achieve musically meaningful harmonic non conformance. To our knowledge, only Markov constraints can compute quickly distributions of solutions of such problems.

Figure 10. Fioritures with various numbers of notes. First one introduces an interesting chromaticism (E to Eb then to D); second example (3 notes) introduce a clearly non conformant chord, that resolves nicely to the D; third example (4 notes) consists in a bold chromatic descent from A minor to D; fourth example (5 notes) uses an interesting triplet-based rhythm that also departs substantially from the A minor chord label; last example is a remarkable jazzy sequence of chords.

Common-sense rhythms

One difficulty that arises when creating fioritures is to find an adequate rhythm for the generated chords. One solution would be to try to imitate rhythm as found in the arranger corpus, but this implies that the corpus used is perfectly reliable, and that metrical information is provided, which is not the case with MIDI files obtained from performances. More importantly, generating Markov sequences with durations raise sparsity issues that do not have general solutions. Another argument is that the rhythm of the fioriture should comply with the genre of the leadsheet more than of the arranger's corpus.

In this study, we have exploited the statistical properties of the leadsheet database to find *commonsense* rhythms that fit with the leadsheet to harmonize. For each rhythm to generate, we query the database to retrieve all the "melodic rhythms" that occur in all jazz standards, at the given metrical position. For a given leadsheet note to harmonize, we retrieve all melodic extracts starting at the same metrical position in the bar, and of the same duration. We then draw a rhythm at random, weighted by its probability in the database. Such a method can be parameterized in many ways (imposing the number of notes, the presence of rests, filter out by composer, genre, etc.). Figure 11 and Figure 12 show the most frequent rhythms found by such a query on the Real book, for 2 different configurations (starting beat in bar and duration).

Figure 11. The 8 most frequent rhythms for a note starting on the first beat of a 4/4 bar with a 4 beat duration, from the Real Book, with their respective frequencies. Query returned 6 062 occurrences of 670 different rhythms.

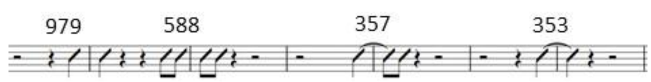

Figure 12. The four most frequent rhythms for a note starting on the last beat of a 4/4 bar with a 2 beat duration, and their respective frequencies. Query found 3943 occurrences of 111 different rhythms.

Full Examples

Two examples of Giant Steps harmonized with fiortiures are given in annex. One in the style of Take 6, and another one in the style of Richard Wagner's tetralogy. In both cases, it can be said that the musical quality is high, compared to previous approaches in automatic harmonization. Preliminary experiments were conducted by playing some harmonizations to highly trained experts (a world famous Brazilian composer, a harmony professor at Goldsmiths College, a talented jazz improviser and teacher, a professional UK jazz pianist): all of them acknowledge that the system produces highly interesting outputs. A full evaluation is under study to try to evaluate precisely the impact of fioritures on the perception of the piece, but is seems reasonable to say that they increase the musical creativity of the software in a significant manner.

Applications to Other Styles

This paper has focused on the style of Take 6, because of the acknowledged difficulty in modeling their productions. Our approach clearly improves on previous attempts at modeling barbershop harmonization such as (Roberts, 2005), who concludes his study by: “although it is possible to formalize the creative process into rules, it does not yield ‘good’ arrangements”. We think we have reached a reasonable level of musical quality here. Our approach, however, is applicable to other styles, as this paper shows with the case of Wagner. Technically our approach is able to harmonize most leadsheets in any style defined by at least one more polyphonic MIDI files, but we did not conduct any specific musical evaluation in other styles yet.

Conclusion

We have introduced the concept of fioriture to harmonize leadsheets in the style of any arranger. Fioritures are controlled random walks within well-defined boundaries defined by long notes in the melody to harmonize. Fioritures could be envisaged under the framework of HMM (as in Farbood and Schoner, 2001). However, HMMs use chord labels as hidden states so we would need an annotated corpus, which is not the case. Furthermore, annotating Take 6 scores with chord labels is in itself an ill-defined problem. Finally, HMM cannot be controlled as precisely and meaningfully as Markov constraints.

Our approach works with non-annotated, non voice-separated corpora for modeling the arranging style. It only requires a definition of chord labels used in the leadsheet (as sets of pitch classes).

Like all music generation systems a rigorous evaluation of our approach is difficult. We claim that our system works remarkably well for most cases, as it rarely makes blatant musical errors, and most often produces musically interesting and challenging outputs. Beyond automatic harmonization, the possibility to control manually fioritures (when, with which parameters) paves the way for a new generation of assisted composition systems. Our approach could be easily extended to exploit social preferences, to help the system choose chords that “sound right” to listeners and ruling out the ones that do not.

Fioritures can also be used as a creative device. By forcing fioritures to have many notes, or by manually substituting chosen leadsheet notes by others, one can generate harmonizations in which the original melody become less and less recognizable, and the style of the arranger becomes increasingly salient. Finally we want to stress that using fioritures to express style is a paradox: fioritures (from italian *fioritura*, flowering) are supposed to be decorative, as opposed to core melody notes, i.e. are not considered primary musical elements. But in our highly constrained context, they can become a device for creative expression.

Acknowledgements

This research is conducted within the Flow Machines project, which received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 291156. We thank A. Dessein for providing us with perfect transcriptions of Take 6 recordings. The accompanying web site[1] gives examples of harmonization of jazz standards.

References

Anders, T. and Miranda, E. R. 2011. Constraint programming systems for modeling music theories and composition. *ACM Comput. Surv.* 43(4): 30.

Ebcioglu, K. 1986. An Expert System for Chorale Harmonization. *Proceedings of AAAI*: 784-788.

Cabral, G. Briot, J.-P., Pachet, F. 2006. Incremental Parsing for Real-Time Accompaniment Systems. *Proc. of 19th FLAIRS Conference*, Melbourne Beach, USA.

Casey, M. A., Veltkamp, R., Goto, M., Leman, M., Rhodes, C., & Slaney, M. 2008. Content-based music information retrieval: Current directions and future challenges. *Proceedings of the IEEE*, *96*(4), 668-696.

Farbood, M. M. and Schoner, B. 2001. Analysis and Synthesis of Palestrina-Style Counterpoint Using Markov Chains. *Proceedings of the 2001 International Computer Music Conference*: 111-117.

Fernández, . D. and Vico, F. J. 2013. AI Methods in Algorithmic Composition: A Comprehensive Survey. *J. Artif. Intell. Res. (JAIR)* 48: 513-582

Koops, H. V. and Magalhães, J. P. and de Haas, W. B. 2013. A functional approach to automatic melody harmonisation. *Proceedings of the first ACM SIGPLAN workshop on Functional art, music, modeling & design (FARM '13),* ACM, New York, NY, USA, 47-58.

Marchini, M. and Purwins, H. 2010. Unsupervised Analysis and Generation of Audio Percussion Sequences. Proceedings of CMMR: 205-218.

Pachet, F. Roy, P. 1995. Integrating constraint satisfaction techniques with complex object structures, 15th Annual Conf. of the British Comp. Society, Cambridge, pp. 11-22.

Pachet, F., Roy, P. and Barbieri, G. 2011. Finite-Length Markov Processes with Constraints. *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*: 635-642, Barcelona.

Pachet, F., Suzda, J., and Martin, D. 2013. A Comprehensive Online Database of Machine-Readable Lead Sheets for Jazz Standards. *Proceedings of ISMIR: 275-280*, Curitiba (Brazil).

Paiement, J.-F. Eck, D. and Bengio, S. 2006. Probabilistic Melodic Harmonization. Proc. of *Canadian Conference on AI: 218-229.*

Roberts, S. 2005. Automated Harmonisation of a Melody into a Four Part Barbershop Arrangement, *Bsc thesis report*, University of Bath.

Steels, L. 1986. Learning the Craft of Musical Composition, Proc. of ICMC.

Take 6. 1988. *Take 6*, Warner Bros. 1988.

Take 6. 1988. *So Much 2 say,* Warner Bros. 1990.

Whorley, R. Rhodes, C. Wiggins, G. and Pearce, M. 2013. Proceedings of the Fourth International Conference on Computational Creativity: 79-86.

[1] http://www.flow-machines.com/harmonization

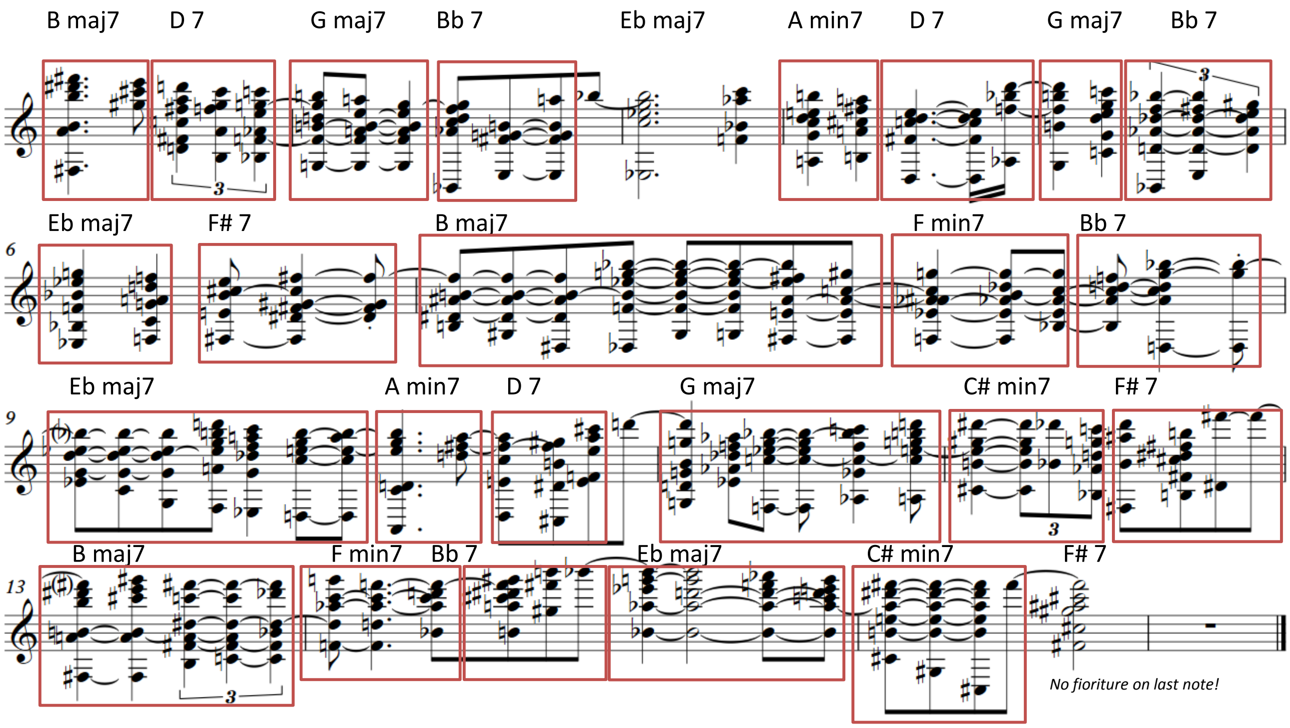

Figure 13. Giant Steps in the style of Take 6 with fioritures of various lengths. Fioritures are indicated by boxes. Note the use of new rhythms and interesting harmonies.

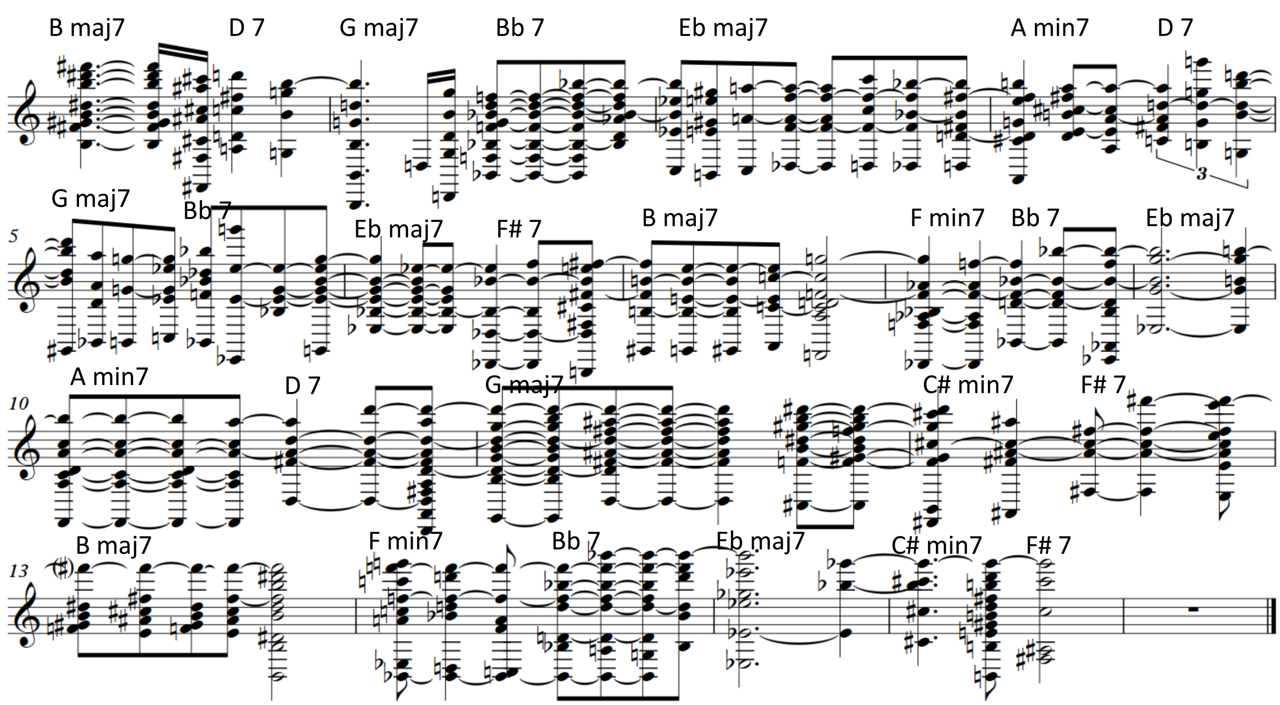

Figure 14. Giant steps with fioritures, in the style of Wagner (training corpus consists of the scores of the Ring tetralogy). The musical output definitely sounds Wagnerian yet follows strictly the Giant Steps leadsheet. Musical comments are available on the accompanying web site.