Revised Selected Papers

Accademia Musicale Studio Musica
Michele Della Ventura, *editor*

2019

Proceedings of the International Conference on New Music Concepts and Inspired Education

Vol. 6

ACCADEMIA MUSICALE
STUDIO MUSICA

Accademia Musicale Studio Musica

International Conference on New Music Concepts and Inspired Education

Proceeding Book
Vol. 6

Accademia Musicale Studio Musica
Michele Della Ventura
Editor

Printed in Italy
First edition: April 2019

www.studiomusicatreviso.it
Accademia Musicale Studio Musica – Treviso (Italy)
ISBN: 978-88-944350-0-9

Preface

This volume of proceedings from the conference provides an opportunity for readers to engage with a selection of refereed papers that were presented during the International Conference on New Music Concepts and Inspired Education. The reader will sample here reports of research on topics ranging from mathematical models in music to pattern recognition in music; symbolic music processing; music synthesis and transformation; learning and conceptual change; teaching strategies; e-learning and innovative learning. This book is meant to be a *textbook* that is suitable for courses at the advanced undergraduate and beginning master level. By mixing theory and practice, the book provides both profound technological knowledge as well as a comprehensive treatment of music processing applications.
The goals of the Conference are to foster international research collaborations in the fields of Music Studies and Education as well as to provide a forum to present current research results in the forms of technical sessions, round table discussions during the conference period in a relax and enjoyable atmosphere.
36 papers from 16 countries were received. All the submissions were reviewed on the basis of their significance, novelty, technical quality, and practical impact. After careful reviews by at least three experts in the relevant areas for each paper, 12 papers from 10 countries were accepted for presentation or poster display at the conference.

I want to take this opportunity to thank all participants who have worked hard to make this conference a success. Thanks are also due to the staff of "Studio Musica" for their help with producing the proceedings. I am also grateful to all members of Organizing Committee, Local Arrangement Committee and Program Committee as well as all participants who have worked hard to make this conference a success.
Finally I want to appreciate all authors for their excellent papers to this conference.

April 2019 — Michele Della Ventura

Contents

Generative Conceptual Blending of High-Level Melodic Features: Shortcomings and Possible Improvements

Maximos Kaliakatsos-Papakostas

Department of Music Studies, Aristotle University of Thessaloniki
maxk@mus.auth.gr

Abstract. Conceptual Blending (CB) theory has been primarily employed as a method for interpreting creative artefacts, while recently it has been used as a creative tool for the algorithmic invention of new concepts. In music, interesting examples have been presented where low-level musical information (e.g. chord roots, chord types or pitch classes) is combined to generate new musical objects (e.g. cadences) or even entire harmonic spaces. These generative frameworks, however, do not incorporate information for high-level descriptive features of music, as pentatonicity or rhythm syncopation. The paper at hand presents a methodology where high-level blending is achieved by recombining low-level information of melodies using a genetic algorithm. A test case is examined where a Chinese Han melody is blended with a Jazz melody and representative blends are analyzed to expose some shortcomings and possible improvements in this approach.

Keywords. Conceptual Blending, Melodic Generation, High-Level Features

1 Introduction

The theory of Conceptual Blending (CB) [1] relates with combinational creativity, which Boden [2] maintains is the hardest to describe formally. CB theory has been used primarily as a method for interpreting creative ideas; such examples can be found in the analysis of Zbikowski [3] on Palestrina's text painting, or the analysis of Tsougras & Stefanou [4] for Mussorgsky's "Pictures at an Exhibition". These analyses mainly examine how musical structures and high-level concepts are blended with extra-musical ideas to form musical works. The methodological framework of CB theory incorporates two input spaces that are described formally, with properties and relations between these properties. The input spaces are blended to form new spaces that include meaningful and creative combinations of properties and relations if the input spaces, potentially leading to the creation of altogether new spaces.

Interesting musical results have been presented lately in computational methodologies where CB is used generatively. For instance, in [5] a methodology was presented where the properties of the Perfect and the Phrygian cadences are combined to generate the tritone substitution cadence, which is a chord progression that was developed in Jazz centuries later in relation to the inputs. A similar approach was employed in the Chameleon[28] melodic harmonization assistant [6], where generic chord transitions, instead of merely cadences, were blended leading to the generation of new probabilistic harmonic spaces that allowed, for instance, meaningful connections between remote tonalities, or the generation of hybrid musical styles, e.g. blends between harmonies of Bach choral and Jazz standards.

The aforementioned interesting results in generative conceptual blending are restricted, however, to low-level descriptions of the musical surface, e.g. root notes of chords, chord types, pitches etc. These low-level descriptions do not represent explicitly high-level concepts, as the ones employed in the interpretative use of conceptual blending. For instance, Zbikowski [3] refers to how Palestrina combines the concept of "falling from heaven" in the lyrics with a "descending" musical passage. The framework for generative conceptual blending used in the "traditional" approach presented in the previous paragraph does not support the incorporation of such high-level features in the blending process.

Aim of this paper is to present a new approach in generative conceptual blending that allows the incorporation of high-level features in the blending process; the specific domain of application is the generation of melodies. With this methodology, high-level concepts of two input melodies can be blended, generating new melodies that reflect the combined high-level information. To this end, Chinese melodies in the Han style and Western melodies, mainly in the style of Jazz standards, have been collected and the results of a specific use case are presented, which expose some shortcomings and ways for possible improvements in this framework. This framework has been examined previously for drums rhythm generation [7], however, the employment of many drums features (forty) did not make clear what the methodology actually does and how it can be improved. In the paper at hand, the focus is placed on four partially distinct features, a fact that allows an accurate assessment of the shortcomings.

2 Methodology for Generating Melodies with High-Level Feature Blending

An overview of the methodology is shown in Fig. 1, according to which, two input melodies are given as input. High-level features are extracted from these melodies and their average first order Markov matrix of pitch transitions is computed. A combination of desired high-level features from the input melodies is selected (input features 1 and 2, corresponding to features from input melodies 1 and 2 respectively); the combined/blended features, along with the average Markov matrix, comprise the target features that a new melody should satisfy. A genetic algorithm is afterwards employed to successively recombine the low-level material of the input melodies until the target feature combination of the input features are achieved. During the evolutionary process, the Markov matrix of pitch transitions for each individual is also computed. The deviation

[28] http://ccm.web.auth.gr/chameleonmain.html

from this matrix from the average Markov matrix of the inputs is introduced as penalty during fitness evaluation (analyzed in Eq. (1) later in the text).

In the middle part of Fig. 1, the typical "blending diamond" is shown, which illustrates the relations between the input spaces, the generic space and the blended space. In this case, the input spaces are sets of four high-level features (numerical values) extracted from the input melodies. The generic space includes features that have similar values; e.g. if two melodies with high rhythm density are given, the generic space is the concept of high rhythm density. The blended spaces is a set of four features, produced by combining features of the input spaces. In the current framework, the generic space is rather trivial since it just indicates which features are of "similar" value. In formal CB theory, however, the generic space plays the crucial role of computing which properties and relations need to be transferred as-they-are in the blend. In the concluding section, the potential role of the generic space in the high-level framework is discussed in the context of future work.

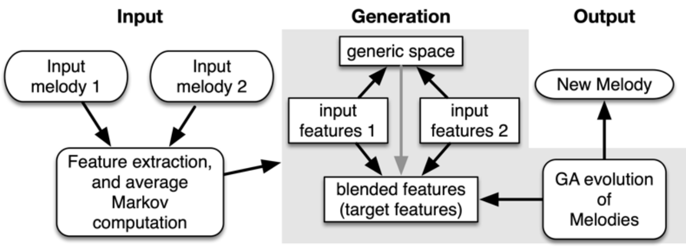

Fig. 1. Overview of the high-level melodic blending methodology. Two melodies are given as inputs and high-level features are extracted. Selected features are blended and used as target features; a genetic algorithm evolves copies of the inputs until a new melody that reflects the target features is generated.

Eq. (1) shows the fitness value for each generated melody, *m*. The α and β constants correspond to the weight of each component during fitness evaluation, respectively deviation from target/blended features and the average Markov matrix of pitch transitions for the inputs. Specifically,

- $f_i(m)$ corresponds to feature *i* of the generated melody,
- f_i^t to target feature *i*,
- $T(m)$ to the Markov matrix of pitch transitions of the new melody and
- $T(I_1, I_2)$ to the average Markov matrix of pitch transitions of the input melodies.

$$f(m) = \alpha \sqrt{\sum_{i=0}^{3} (f_i(m) - f_i^t)^2 + \beta \, \|T(m) - T(I_1, I_2)\|} \qquad (1)$$

The number of employed high-level features, as shown in the sum of feature differences in Eq. (1), is four. These features are the following:

1. *Rhythm density*: measured as the ratio of the number of note events in the melody over the total number of sixteenths (maximum considered resolution) in the melody.
2. *Syncopation*: a number that reveals how complicated a rhythm is, computed with a simplified version of the syncopation version proposed in [8].
3. *Pentatonicity*: this feature describes the best correlation between the pitch class profile of the melody and any circular shift of a binary pentatonic template. In essence, this feature produces a value between 0 and 1 that shows how close to the pentatonic a pitch class profile distribution is.
4. *Small intervals*: the percentage of non-zero intervals between successive pitches that are less or equal to two semitones.

The evolutionary process is inspired by low-level blending. The main idea is to recombine only elements of the input melodies, towards approaching the target features. This means that genetic operators related to mutation are excluded. This also means that the initial population comprises solely copies of the two input melodies, no randomly selected melodies or random melodies altogether. The genetic operators act on two parent melodies selected each time from the pool of melodies in the current generation allowing for: (a) random exchange of a single pitch; (b) exchange of the rhythmic structure in randomly selected bars with preservation of pitches; and (c) exchange of the entire rhythm structure in the entire melodies, with preservation of pitches. These operators ensure that only pitches and rhythm structures in the input melodies is recombined for forming new individuals/melodies towards capturing the target/blended features.

3 An Analysis of Generated Blends Towards Identifying and Improving Shortcomings

Multiple melodies can be generated by applying the methodology described previously. Specifically, multiple blends can be created by combining high-level features from each input and many melodies can be generated that correspond to these feature combinations. All combinations (or blends) of the four high-level features that are employed in this experimental setup are fourteen in number; actually, all combinations are sixteen but the two inputs are included as the trivial "combinations" (taking all features from each input does not lead to a blend).

Two stylistically distinct sets of melodies were gathered: 350 melodies from the Chinese Han style, taken randomly from the over 1200 melodies in the Essen Corpus [9-10], and 350 Western melodies (mainly Jazz standard melodies) obtained from online resources. An illustration of all the melodies, along with some blends of specific input melodies, is given in Fig. 2. Therein, Principal Component Analysis (PCA) is applied on the four-dimensional representation of all melodies and the projection of the first two dimensions is shown. The first two PCA dimensions account for the 72% of data variance. The horizontal axis of the PCA plot corresponds to a combination of the pentatonicity an density features. Specifically, the horizontal axis has a negative correlation of -0.83 with the density feature and -0.90 with pentatonicity, while (absolute) correlation increases even more for the sum of these two features to -0.98. This means that the horizontal axis can be interpreted as follows: melodies to the left have higher pentatonicity and density feature values, while melodies to the right have lower values for these features. Therefore, as Fig. 2 shows, Han melodies have higher pentatonicity and density values and Western melodies the opposite. The vertical axis is correlated with the syncopation feature (0.87), while there is no clear distinction between Han and Western melodies regarding this feature – both styles are similarly represented in all syncopation ranges. The small intervals feature does not appear to be correlated with any of the PCA axes.

During numerous sessions, many pairs of melodies from the two styles have been selected for blending; multiple blending results were examined and a representative case is used as an example. Fig. 2 shows the pair of inputs used in this example and with the resulting blends, on the two-dimensional PCA projection. The two inputs incorporate distinctly different characteristics, as indicated by the feature values in the first two columns in Table I; the scores of the two melodies are shown in Fig. 2 (a) and (b), for the Han and the Jazz inputs respectively. By Fig. 2 one can identify four clusters of blends: two clusters forming the two respective inputs and two clusters in the center, one higher and one lower in the vertical dimension. Some representative of each cluster are shown in Table I (their feature values) and Fig. 2 (c)-(f).

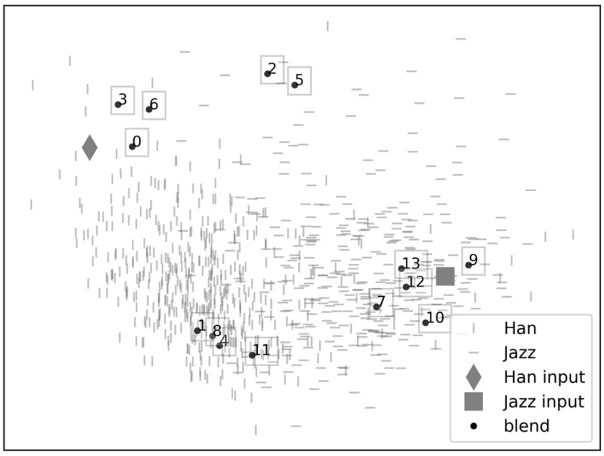

Fig. 2. Illustration of the first two PCA dimensions for a collection of Han and Jazz melodies, along with two inputs (Han and Jazz) and all the best blends for each feature combination.

Table I, specifically, shows features related to the Han input with bold numbers. For instance Blend 6 has syncopation and pentatonicity values coming from the Han input (0.66 and 0.99 respectively) and rhythm density and small intervals from the Jazz input (0.37

and 0.74). The genetic algorithm that implements this blend, however, fails to capture the rhythm density feature according to the target value of the Jazz input (0.26) and remains relatively high in the blend (0.37). On the other hand, blend 12 has low rhythm density and pentatonicity values, taken from the Jazz input, and thus it is placed on the right side of the graph, closer to the Jazz input. Again the genetic algorithm fails to assign the high targeted value for syncopation (0.63) from the Han input and compromises with a lower value (0.30).

Shortcomings: Blends 6 and 12 make it obvious that there can be cases where the genetic algorithm cannot find a good compromise between potentially conflicting feature values and generates melodies that are unavoidingly not optimal in terms of one such feature. Rhythm density and syncopation seems to be a good example of that. Even though the notions of rhythm density and syncopation are not necessarily related in general, in the examined example and for the selected inputs higher density means higher syncopation (Han) and vice versa (Jazz). Therefore, the genetic algorithm does not have the necessary genetic material, nor the genetic operators, to construct blends that successfully violate the apparent correlation.

TABLE I: FEATURES OF THE INPUTS AND SELECTED BLENDS.

Melody	Rhythm Density	Syncopation	Pentatonicity	Small Intervals
Han	**0.50**	**0.63**	**0.99**	**0.43**
Jazz	0.26	0.00	0.36	0.76
Blend 6	0.37	**0.66**	**0.99**	0.74
Blend 12	0.19	**0.30**	0.44	0.79
Blend 11	0.23	0.00	**0.99**	0.75
Blend 2	**0.50**	**0.64**	0.43	**0.44**

Having a pool of many features to blend for generating a new melody increases the chances that some feature combinations become potentially contradicting. On the other hand, using very few features might lead to an under-definition of desired properties, leading to melodies that are typically correct in terms of features, but do not capture the overall essence of what the features aim to reflect. Blend 2 in Fig. 3 (f) is an example of that: it accurately captures the features the target features and pentatonicity and small intervals are kept low (0.43 and 0.44 respectively) but in a peculiar manner that does not reflect the intended properties according to the given inputs. Low pentatonicity should make the pitch content of the blend similar to the one of the Jazz input, hence presenting rich pitch content. Contrarily, the result (Fig. (f)) merely includes three pitches, while mainly one pitch is played; this pitch setup satisfies the low pentatonicity demand but the pitch content does not reflect the characteristics of the low pentatonicity (Jazz) input. Similarly, a low percentage of small intervals is achieved by playing mostly constant intervals, which are not accounted for when computing the small intervals feature as described in Section 2.

Fig. 3. Han (a) and Jazz (b) inputs, along with selected blends, namely Blend 6 (c), Blend 12 (d), Blend 11 (e) and Blend 2 (f).

Possible improvements: Incompatibility of certain feature combinations could possibly be addressed by applying multi-objective optimization criteria, allowing the genetic algorithm to find a good compromise in violating, to the least possible extent, a combination of correlated feature values instead of completely ignoring the value of one feature. Additionally, the process could possibly allow the inclusion of "genetic material" from other melodies (not necessarily the inputs) that potentially satisfy the required feature combinations. Keeping material (pitches and rhythms) solely from the inputs, to some extent reflects the basic principles of low-level blending. The introduction of new material is, however, "permitted" in the formal framework of blending, according to the notion of "blending *completion*". Blending completion is the process of incorporating elements that do not exist in the inputs, in order to resolve incomplete or inconsistent blends. In the example of the tritone substitution cadence example, the introduction of the Ab pitch in the penultimate chord is a product of blending completion, which ensures that the type of the penultimate chord is major with minor seventh.

The under-definition of features could be addressed by involving implicit machine learning techniques in the process of evaluation that guarantee conformation of the blended melody with some critical stylistic (implicitly learned and latent) features. It is reminded that under-definition of a feature may lead to the generation of melodies that satisfy this specific feature, but skip to satisfy some other critical aspects of the style (e.g. low pentatonicity is achieved with unusually extensive note repetitions in the case of Blend 2). Trying to over-define a feature by introducing more related features (e.g. combining pentatonicity with the feature of note repetitions or with pitch class profile information entropy) will possibly lead to conflicting relations between them, amplifying the problem of feature incompatibility described in the previous paragraph. Implicit learning methods learn latent features that are not necessarily interpretable by iterating through data; an example of such class of methods is deep neural networks.

It appears that such techniques, having no explicit information about what pentatonicity or repetition is, can play the role of a generic evaluator that is not biased towards any specific feature. For instance, imagine a neural network, *N*, that has been trained on many melodies from many styles. When evaluating Blend 2, such an evaluator would reject this melody because it would be "unusual" (regardless of the fact the reason might not be interpretable) in comparison to the other melodies that the evaluator has been trained on. Eq. (2) shows how fitness evaluation in Eq. (1) could change in order to accommodate the probability given *N* for the melody *m*; this probability is shown in Eq. (2) as *N(m)*.

$$f(m) = \alpha \sqrt{\sum_{i=0}^{3}(f_i(m) - f_i^t)^2} + \beta \; \|T(m) - T(I_1, I_2)\| + \gamma \; N(m) \qquad (2)$$

4 Conclusion

This paper has presented a methodology that allows the incorporation of high-level features in generative conceptual blending. While current approaches allow only blending of low-level information, the presented approach is takes recombines low-level information towards achieving targets that are related to higher-level descriptive values. This methodology is applied on blending melodies in the Chinese Han and Western, mainly Jazz, styles. The algorithmic core is based on conceptual blending, while the generative part introduces a genetic algorithm that recombines material from the input melodies. Four distinct high-level features were introduced, which is a number of features small enough to allow qualitative assessment of shortcomings and possible improvements in the examined framework. An example test-case has been presented, where a Han melody is blended with a Jazz melody and representative samples of fourteen blends are examined in detail.

Two main shortcomings were identified that, under specific circumstances, might lead to the generation of sub-optimal blended melodies: (a) the inclusion of conflicting features in the blend and (b) the potentially low descriptive quality of a feature. Conflicting features could be addressed by introducing multi-objective optimization criteria in the evolutionary process, or/and by allowing material from third melodies (not solely the inputs) to be involved in the evolutionary stage. Low descriptive quality of features could be addressed by introducing implicit learning models (e.g. deep learning) that play the role of stylistic evaluators, ensuring that the blended melodies incorporate some critical stylistic properties. Future work will also incorporate a more active role of the generic space in this high-level framework. In this version of the feature blending algorithm, the generic space simply indicates which features of the inputs are similar. In a future version, the inclusion of low-level "similar" patterns of the inputs will also be included. Thereby, similar segments/patterns in the input will be identified and included in the generic space and, subsequently, in all generated blends.

Acknowledgements

This work has been supported by the Greek National Scholarships Foundation (Grant Agreement Number: 2016-050-0503-8138).

References

[1] G. Fauconnier, M. Turner, The Way We Think: Conceptual Blending And The Mind's Hidden Complexities, reprint edition Edition, Basic Books, New York, 2003.

[2] M. A. Boden, Computer Models of Creativity, AI Magazine 30 (3) (2009) 23.

[3] L. M. Zbikowski, Conceptualizing Music: Cognitive Structure, Theory, and Analysis, Oxford University Press, 2002.

[4] C. Tsougras, and D. Stefanou, Embedded blends and meaning construction in Modest Musorgsky's Pictures at an Exhibition, Musicae Scientiae 22.1 (2018): 38-56.

[5] M. Eppe, R. Confalonier, E. Maclean, M. Kaliakatsos-Papakostas, E. Cambouropoulos, M. Schorlemmer, M. Codescu, K.-U. Ku ̈hnberger, Computational invention of cadences and chord progressions by conceptual chord-blending., in: International Joint Conference on Artificial Intelligence (IJCAI) 2015, 2015.

[6] M. Kaliakatsos-Papakostas, M. Queiroz, C. Tsougras, E. Cambouropoulos, Conceptual blending of harmonic spaces for creative melodic harmonisation, Journal of New Music Research 43 (3) (2017) 1–25.

[7] M. Kaliakatsos-Papakostas, Generating drum rhythms through data- driven conceptual blending of features and genetic algorithms, in: Evolutionary and Biologically Inspired Music, Sound, Art and Design, Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2018, pp. 145–160.

[8] G. Sioros, M. Miron, M. Davies, F. Gouyon, G. Madison, Syncopation creates the sensation of groove in synthesized music examples, Frontiers in psychology 5, 2014.

[9] H. Schaffrath, Beyond midi, MIT Press, Cambridge, MA, USA, 1997, Ch. The Essen Associative Code: A Code for Folksong Analysis, pp. 343–361.

[10] H. Schaffrath, D. Huron, The Essen folksong collection in kern format.[computer database], Menlo Park, CA: Center for Computer Assisted Research in the Humanities.

This book presents a collection of selected papers that present the current variety of all aspect of music research, development and education, at a high level. The respective chapters address a diverse range of theoretical, empirical and practical aspects underpinning the music science and teaching and learning, as well as their pedagogical implications. The book meets the growing demand of practitioners, researchers, scientists, educators and students for a comprehensive introduction to key topics in these fields. The volume focuses on easy-to-understand examples and a guide to additional literature.

Michele Della Ventura, editor

New Music Concepts and Inspired Education

Revised Selected Papers

ISBN: 978-88-944350-0-9

9 788894 435009

www.studiomusicatreviso.it