

An Argument-based Creative Assistant for Harmonic Blending

Maximos Kaliakatsos-Papakostas^a, Roberto Confalonieri^b, Joseph Corneli^c,
Asterios Zacharakis^a and Emiliós Cambouropoulos^a

^aDepartment of Music, Aristotle University of Thessaloniki, Greece

^bArtificial Intelligence Research Institute, IIIA-CSIC, Bellaterra (Barcelona), Spain

^cDepartment of Computing, Goldsmiths, University of London, UK

^a{max,aszachar,emilios}@mus.auth.gr

^bconfalonieri@iia.csic.es

^cj.corneli@gold.ac.uk

Abstract

Conceptual blending is a powerful tool for computational creativity where, for example, the properties of two harmonic spaces may be combined in a consistent manner to produce a novel harmonic space. However, deciding about the importance of property features in the input spaces and evaluating the results of conceptual blending is a nontrivial task. In the specific case of musical harmony, defining the salient features of chord transitions and evaluating invented harmonic spaces requires deep musicological background knowledge. In this paper, we propose a creative tool that helps musicologists to evaluate and to enhance harmonic innovation. This tool allows a music expert to specify arguments over given transition properties. These arguments are then considered by the system when defining combinations of features in an idiom-blending process. A music expert can assess whether the new harmonic idiom makes musicological sense and re-adjust the arguments (selection of features) to explore alternative blends that can potentially produce better harmonic spaces. We conclude with a discussion of future work that would further automate the harmonisation process.

Introduction

The invention of new harmonic spaces in this paper is conceived as a computational creative process according to which a new harmonic idiom is created by means of blending the ‘atoms’ of harmony, i.e., transitions between chords. The blended transitions are created by combining the features characterising pairs of transitions belonging to two idioms (expressed as sets of potentially learned transitions) according to an amalgam-based algorithm (Confalonieri et al., 2015a; Eppe et al., 2015b) that implements Fauconnier and Turner (2002)’s theory of conceptual blending. The transitions are then used in an extended harmonic space that accommodates the two initial harmonic spaces, linked with the new blended transitions.

When modeling creative processes computationally, one of the key questions is how good are the created artefacts. The approach to evaluation that has been applied most frequently within computational creativity requires a human to evaluate attributes of the created work or the system’s operation. Basic measures consider the *typicality* of a generated artefact within a particular genre, or the *quality* of the generated work according to the users’ aesthetic judgement (Ritchie, 2007).

In music blending, the evaluation of artefacts is not a trivial matter. This is due not only to the time evolving nature of the final output, but also to the lack of clearly defined criteria for their assessment. In the particular case of transition¹ blending, which is how harmonic blending is approached in this paper, the evaluation of the blends is of key importance, in order to produce musically meaningful extended harmonic spaces. To evaluate the set of blended transitions and the corresponding generated extended harmonic space, several musical features need to be taken into account according to indications by musicologists. The importance of each particular feature, however, is not known in advance and musicologists need to make adjustments by experimenting with a large set of test cases.

To ease this task, in this paper, we propose a creative tool (Figure 1) that assists a musicologist with the evaluation of harmonic blends. The system allows a musicologist to specify *arguments*—abstracting the properties of chords and transitions—and to use them for an iterative evaluation of the blended outcome, based on the transitions that the system proposes in order to connect two (potentially remote) harmonic spaces.

Using arguments to make and explain decisions has been proposed and explored in Artificial Intelligence (Bench-Capon and Dunne, 2007), where an argument is a reason for believing a statement, choosing an option, or doing an action. In most existing works on argumentation, an argument is either considered as an abstract entity whose origin and structure are not defined (Dung, 1995), or it is a logical proof for a statement where the proof is built from a knowledge base (Amgoud and Prade, 2009). The use of argumentation in concept invention is, on the other hand, less frequent. Confalonieri et al. (2015b) use Lakatosian’s reasoning to model dialogues in which users engage to discuss the intended meaning of blended concept.

In our approach, arguments encapsulate desirable properties that the user would like to have in the resulting transition blends. Arguments are specified by the user by answering specific questions over the features of the idioms selected as input for the transition blending process. Providing some

¹For the rest of the paper, the term transition will be referring to a pair of chords where one follows the other. For example, G7 → C is a transition describing the perfect cadences in tonal music.




Figure 1: A schematic diagram of the system’s workflow.

higher level arguments as inputs to the system is equivalent to allowing a musical expert to interact with it in a language he/she understands. This offers the user a flexible way to adjust the harmonic blending properties according to different input scenarios in order to improve the creativity of the system. Extended experimentation with the system—by making use of the available arguments—can enable music experts to provide valuable feedback regarding the functionality of transition features, thus directly intervening in the blending process by answering simple questions. In a future scenario, the assessment of the system will be based on merely musicological criteria that should be more clearly defined.

The paper is organised as follow. In the next section, we describe the harmonic blending creative process embedded in a creative assistant tool we implemented. Next, we describe the methodology of transition blending and extending harmonic spaces. We show how user arguments are used to evaluate transition blends based on two criteria resembling two of the optimality principles of conceptual blending. Then, we present a process-based system evaluation that focuses on the creative acts of programmers (Colton et al., 2014). This evaluation is helpful in guiding further developments of the system. These are discussed in a concluding section.

System overview and test cases

Figure 1 illustrates a diagram of the presented system. The user (music expert) interacts with the system through the Graphical User Interface (GUI), where she/he selects two initial idioms (harmonic spaces expressed as sets of permitted chord transitions) in γ and defines the important features used in conceptual blending by answering specific questions (argumentation) in β . The selected initial idioms are described as sets of chord transitions, while the provided answers to questions are mapped to enabling/disabling features of transitions (see Section ‘Chord transitions description and blending’) that define the outcome of transition

blending (see Section ‘Evaluation of transition blending via arguments’).

Afterwards, pairs of transition in the two initial harmonic spaces are given as inputs to the transition blending system in ϵ where new transitions are invented through conceptual blending.² These transitions are then integrated into an *extended* musical idiom that includes the initial idioms selected by the user. While the role of the new transitions is to provide musically meaningful connections between the initial harmonic spaces. The created extended idiom is displayed to the user in the GUI in terms of a transition matrix (see Section ‘From transition blends to transition matrices’). By observing the matrix, the music expert evaluates (ι) the results produced by the current blending setup, i.e., the given questions to the argumentation module (β), and re-adjusts her/his answers in β accordingly.

Several scenarios for initial idiom combinations are available to the user. The system includes several harmonic blending test cases according to which the user can blend simple ‘artificial’ harmonic spaces as well as harmonic spaces trained from data in different tonalities. The artificial harmonic spaces are manually constructed to include simple transitions that can typically be found in tonal music in order to allow clear interpretations of the results, e.g., a C major space included the chords C, F and G7. Among the trained idioms that have been examined, there are sets of Bach chorales in major and minor mode, and sets of modal chorales in several modes.

The test cases, in which harmonic spaces in different tonalities are blended, resemble the musical task of finding transition paths for tonality modulations (changing the tonality of a given harmonic space). This task allowed music experts to identify arguments for defining the features of transition blending that connect potentially remote harmonic spaces (e.g., C major with F \sharp major) in a manner that is explainable in music theory in terms of tonality modulations. Through the processes offered by the system, the music experts were able to come to conclusions about what transition features are important for constructing meaningful connections between different combinations of pairs of initial harmonic spaces.

Methodological aspects of transition blending and extending harmonic spaces

The cognitive theory of conceptual blending by Fauconnier and Turner (2002) has been extensively used in linguistics, music composition (Zbikowski, 2002), music cognition (Antovic, 2009, 2011) and other domains mainly as an analytical tool, which is useful for explaining the cognitive process that humans undergo when engaged in creative acts. According to this theory, human creativity is modeled

²For instance, if the two initial harmonic spaces are a tonal C minor (I_1) and a modal C phrygian space (I_2), then a pair of transitions for blending could be $G7 \rightarrow Cm$ (from I_1) and $B\flat m \rightarrow Cm$ (from I_2). A possible resulting transition from blending these input transitions is the tritone substitution cadence, $C\sharp 7 \rightarrow Cm$, as computed in Eppe et al. (2015b) and Zacharakis, Kaliakatsos-Papakostas, and Cambouropoulos (2015).




Figure 2: Conceptual blending based on amalgam. The generic space is computed (1) and the input spaces are successively generalised (2), while new blends are constantly created (3). Some blends might be inconsistent or purely evaluated according to blending optimality principles or domain specific criteria.

as a process by which a new concept is constructed by taking the commonalities among two *input spaces* into account, to form a so-called *generic space*, and by projecting their non-common structure in a selective way to a novel blended space, called a *blend*.

In computational creativity, conceptual blending has been modeled by Goguen (2006) as a generative mechanism, according to which input spaces are modeled as *algebraic specifications* and a blend is computed as a categorical *colimit*. A computational framework that extends Goguen’s approach has been developed in the context of the COncEpt INVENTion Theory³ (COINVENT) project (Schorlemmer et al., 2014) based on the notion of *amalgams* (Ontañón and Plaza, 2010). According to this framework, *input spaces* are described as sets of features, properties and relations, and an *amalgam*-based workflow finds the blends (Confalonieri et al., 2015a; Eppe et al., 2015b). The amalgam-based workflow generalises input concepts until a generic space is found and ‘combines’ generalised versions of the input spaces to create blends that are consistent or satisfy certain properties that relate to the knowledge domain (Figure 2).⁴

Amalgam-based conceptual blending has been applied to invent chord cadences (Eppe et al., 2015a; Zacharakis, Kaliakatos-Papakostas, and Cambouropoulos, 2015). In this setting, cadences are considered as special cases of chord transitions—pairs of chords, where the first chord is followed by the second one—that are described by means of features such as the roots or types of the involved chords, or intervals between voice motions, among others. When blending two transitions, the amalgam-based algorithm first finds a generic space between them (point 1 in Figure 2). For instance, in the case of blending the perfect with the

³<http://www.coinvent-project.eu>

⁴In the process of blending through amalgams, the notions of ‘amalgam’ and ‘blend’ are the same. Therefore, in the following paragraphs they are used interchangeably.




Figure 3: Example of blending cadences, which are special case of transitions, where blending the perfect and the Phrygian produce the tritone substitution cadence blend.

Phrygian cadences (Figure 3)—described by the transitions $I_1: G7 \rightarrow C$ and $I_2: B\flat m \rightarrow C5$ respectively—their generic space consists of any transition that has a second chord with the root note C, since this is the root note of both inputs’ second chords (C and C5).

After a generic space is found, the amalgam-based process computes the amalgam of two input spaces by *unifying* their content. If the resulting amalgam is inconsistent, then it iteratively generalises the properties of the inputs (point 2 in Figure 2), until the resulting unification is consistent (point 3 in Figure 2). For instance, trying to directly unify the transitions $I_1: G7 \rightarrow C$ and $I_2: B\flat m \rightarrow C5$ would yield an inconsistent amalgam, since a transition cannot both include and *not* include a leading note to the second chord’s tonic (which is a property of I_1 and the I_2 respectively). Therefore, the amalgam-based process generalises the clashing property in one of the inputs (e.g., the property describing the absence of leading note would be left empty in I_2) and tries to unify the generalised versions of the inputs again. After a number of generalisation steps are applied (point 2 in Figure 2), the resulting blend is consistent (point 3 in Figure 2). However, it may be the case that the blend is not complete, in the sense that this process may have generated an over-generalised term.

After several blends have been computed, an automated evaluation process ranks them according to some optimality principles (Fauconnier and Turner, 2002, Chapter 16). These principles are a necessary aspect of conceptual blending since they allow to filter interesting blends from the (potentially too) many possible ones.⁵ A complete description of optimality principles is out of the scope of this paper and the reader is referred to Goguen and Harrell (2010) for applications of such principles in the *Alloy* algorithm. We give, however, two extreme examples of ‘bad blends’ for clarifying the importance of using optimality principles in conceptual blending. *Example 1*: Each of the input spaces is a trivial form of a (bad) blend, since no information from the other input spaces is considered. *Example 2*: A blend that includes all properties of the generic space, but no other information of the inputs spaces is a bad blend, since it has the least possible connections with the input spaces. These examples suggest two criteria for ranking the blends; we provide a computational characterisation of them below.

⁵The amalgam-based algorithm produces many blends by following alternative generalisation paths.

Chord transitions description and blending

Individual chord transitions are the ‘atoms’ of the methodology followed herein to construct new transition matrices. Specifically, transition sets from two musical idioms provide input transitions for blending, producing a list of blended transitions that are afterwards embedded in an extended harmonic space. This methodology is described briefly in the next section while some definitions regarding chord transitions follow.

Definition 1. A chord transition c is described by a set of features \mathcal{F} .

In this work a transition is represented by 17 features. Features 1-6 refer to the involved chords. Features 8 to 10 indicate changes during the transitions and are based on the Directed Interval Class (DIC) vector (Cambouropoulos, Katsivalos, and Tsougras, 2013; Cambouropoulos, 2012). Feature 7 accounts for the change that occurred regarding the chords’ root notes. The features considered important in this work are the following:

1. *fromRoot*: the root pitch class of the first chord,
2. *toRoot*: the root pitch class of the second chord,
3. *fromType*: the type of the first chord (GCT base),
4. *toType*: the type of the second chord (GCT base),
5. *fromPCs*: the pitch classes included in the first chord,
6. *toPCs*: the pitch classes included in the second chord,
7. *DICinfo*: the DIC vector of the transition,
8. *DIChas0*: Boolean value indicating whether the DIC of the transition has 0,
9. *DIChas1*: As above but for DIC value 1,
10. *DIChasMinus1*: As above but for DIC value -1 ,
11. *ascSemZero*: Boolean value indicating whether the first chord has the relative pitch class value 11,
12. *descSemZero*: As above but value 1,
13. *semZero*: As above but for value 11 or 1,
14. *ascSemNextRoot*: Boolean value indicating whether the first chord has a pitch class with ascending semitone relation with the pitch class of the second chord’s root,
15. *descSemNextRoot*: As above but with descending semitone,
16. *semNextRoot*: As above but with either ascending or descending semitone, and
17. *5thRootRelation*: Boolean value indicating whether the first chord’s root note is a fifth above of the second’s.

Each feature can be considered as a function that assigns a value to a chord transition c . Features’ values are defined differently depending on the properties they represent. For instance, features 3 to 8 are set-value functions that assign a set of values to a chord. We refer to them as $F_i(c)$. The value of the feature 7 is a vector and we refer to it as $\vec{f}(c)$. Finally, all the other features are binary functions and we refer to them as $f_i(c)$.

From transition blends to transition matrices

In the literature, an effective and common way to describe chord progressions in a music idiom in a statistical manner is by using first-order Markov models (see Kaliakatsos-Papakostas and Cambouropoulos (2014); Simon, Morris, and Basu (2008), among others). Such models reflect the probabilities of each chord following other chords in the idiom, as trained or statistically measured throughout all the pieces in the examined idiom. In this context, individual transitions play an important role on indicating particular characteristics of an idiom.

A convenient way to represent a first order Markov model is through transition matrices, which include one respective row and column for each chord in the examined idiom. The probability value in the i -th row and the j -th column exhibits the probability of the i -th chord going to the j -th —the probabilities of each row sum to unit. The utilised chords are actually represented by chord group exemplars, obtained by the method described in Kaliakatsos-Papakostas et al. (2015), while transitions between chords that pertain to the same chord group are disregarded. The representation of chords is based on the General Chord Type representation (Cambouropoulos, Kaliakatsos-Papakostas, and Tsougras, 2014).

Then, an important question is: *How would a blended idiom be expressed in terms of a transition matrix, provided that the transition matrices of two initial idioms are available?*

The idea examined in the present system is to create an *extended* transition matrix that includes new transitions that allow moving across chords of the initial idioms by potentially using new chords. The examined methodology uses transition blending to create new transitions that: (a) maximally preserve the common parts of transitions between the two initial spaces, and (b) incorporate blended characteristics for creating a smooth ‘morphing’ harmonic effect when moving from chords of one space to chords of the other. An abstract illustration of an extended matrix is given in Figure 4.

By analysing the graphical representation of an extended matrix as depicted in Figure 4 the following facts are highlighted:

1. By using transitions in I_i , only chords of the i -th idiom are used. When using these transitions, the resulting harmonisations preserve the character of idiom i .
2. Transitions in A_{i-j} create direct jumps from chords of the i -th to chords of the j -th idiom. Blended transitions in A_{i-j} can be directly included in the extended matrix.
3. Transitions in B_{i-x} constitute harmonic motions from a chord of idiom i to a newly created chord by blending. Similarly, transitions in the B_{x-j} arrive at chords in idiom j from new chords. For moving from idiom i to idiom j using one external chord c_x that was produced by blending, a chain of two transitions is needed: $c_i \rightarrow c_x$ followed by a transition $c_x \rightarrow c_j$, where c_i in idiom i and c_j in idiom j respectively. A chain of two consecutive transitions with one intermediate external chord from chords of i to chords of j will be denoted as B_{i-x-j} .




Figure 4: Graphical description of an *extended* matrix that includes transition probabilities of both initial idioms and of several new transitions generated through transition blending. These new transitions allow moving across the initial idioms, creating a new extended idiom that is a superset of the initial ones.

- Sector C transitions are disregarded since they incorporate pairs of chords that exist outside the i -th and j -th idioms, violating our hypothesis for moving from one known chord sets to the other using one new chord at most.

Based on this analysis of the extended matrix, a methodology is proposed for using blends between transitions in I_1 and I_2 . Thereby, transitions in I_1 are blended with those in I_2 and a certain number of best blends is stored for further investigation, creating a pool of best blends. Based on multiple simulations, a large number of the best blends (i.e. 100) in each blending simulation should be inserted in the pool of best blends (\mathcal{B}), so that several commuting scenarios can be created between the initial transition spaces. Thus, a greater number of blends in the pool of best blends introduces a larger number of possible commuting paths in A_{i-j} or in B_{i-x-j} .

Evaluation of transition blending via arguments

By applying the aforementioned blending process a pool of best blends is created that is afterwards used for connecting the transition blocks of two initial idioms, through forming an extended matrix. When a music expert is using the system, she/he is able to select pairs of initial input idioms⁶, choose which aspects of blending are important through arguments (analysed in the following paragraphs) and evaluate/re-adjust this choice by observing the produced results in the extended matrix.

The user evaluates the importance of several transition features by answering questions based on the connecting transitions produced by blending in the extended matrix. The features related to the transitions and their constituent chords are classified into 9 questions (Table 1).

⁶Except from the idioms used in this paper, new idioms can (a) be manually constructed by providing permitted chord transitions and their probabilities, or (b) learned from data using the first order Markov matrix of chord transitions.

Question	Chord Properties	Transition Changes
Q1	<i>fromRoot</i>	
	<i>toRoot</i>	
	<i>fromType</i> <i>toType</i>	
Q2	<i>fromRelPCs</i> <i>toRelPCs</i>	
Q3		<i>DIChas0</i>
Q4		<i>DIChas1</i> <i>DIChasN1</i>
Q5		<i>DIChas2</i> <i>DIChasN2</i>
Q6		<i>DICinfo</i>
Q7	<i>ascSemZero</i>	
	<i>descSemZero</i>	
	<i>semZero</i>	
Q8	<i>ascSemNextRoot</i>	
	<i>descSemNextRoot</i>	
	<i>semNextRoot</i>	
Q9		<i>5thRootRel</i>

Table 1: Abstraction of chords' and transition changes' features.

- Q1:** Are roots and types of chords important?
Q2: Are individual pitch classes of chords important?
Q3: Are repeating pitch classes in transitions important?
Q4: Are semitone steps in transitions important?
Q5: Are tone steps in transitions important?
Q6: Are the intervallic contents of transitions important?
Q7: Are semitone motions to the tonic important?
Q8: Are semitones to the second chord's root important?
Q9: Are motions of the chord roots by 5th important?

The first two questions concern characteristics of the chords that constitute the transition (features 1 to 6), while the remaining seven questions concern intervallic changes that occur within the transition (features 7 to 17). Relating questions to transition features was performed with the involvement of music experts, to ensure that the mapping is as accurate and as informative to the user as possible.

We denote the set of questions available to the user as \mathcal{Q} . When a user selects a question, an argument is automatically generated. For the sake of this paper, an argument is defined as follows.

Definition 2. An argument A is a tuple $\langle q, F \rangle$, where $q \in \mathcal{Q}$ and $F \subset \mathcal{F}$.

The user can specify at most 9 arguments, each of them is mapped to a set of properties. The set of user arguments $\{A_1, \dots, A_9\}$ corresponding to answers to \mathcal{Q} will be denoted by \mathcal{A} . We assume to have a function $\psi : \mathcal{A} \rightarrow \mathcal{F}$ that returns the set of chord and transition properties associated with an argument (e.g., for the purposes of the current analysis, Table 1 specifies ψ as a look-up function). The arguments are used to compute two criteria in order to rate a blend: *total association* and *symmetry*.

The total association indicates the total number of properties that a blend inherits from the inputs. A blend with higher input associations is preferable since it is structurally more deeply related with the inputs. The total association is calculated by taking the individual association of a blend w.r.t. the input chord transitions into account. The individual association of a blend b w.r.t. to an input I , denoted as $a(b, I)$, is defined as:

$$a(b, I) = \sum_{A_i \in \mathcal{A}} \text{Val}(A_i, b, I)$$

where $\text{Val} : \mathcal{A} \rightarrow \mathbb{R}$ is a function that takes an argument as input and aggregates the values of the chord and transition change properties associated with the argument, by interpreting them according to some music background knowledge. Depending on the type of argument, Val is defined in different ways.

When an argument refers to the roots and types of chords (A_1), Val is defined as:

$$\text{Val}(A_1, b, I) = \sum_{F_j \in \psi(A_1)} \text{equals}(F_j(I), F_j(b))$$

The value of A_1 is calculated by counting how many properties—among *fromRoot*, *toRoot*, *fromType* and *toType*—are equals between a blend b and an input I . *equals* is a function that returns 1 when two sets are equals and 0 otherwise.

When an argument refers to the individual pitch classes of chords (A_2), Val is defined as:

$$\text{Val}(A_2, b, I) = \sum_{F_j \in \psi(A_2)} |F_j(I) \cap F_j(b)|$$

The value of A_2 is calculated as the number of elements that are common in the set-value properties *fromRelPCs* and *toRelPCs* of a blend b and an input I .

When an argument refers to the intervalic contents of transitions (A_6), Val is defined as:

$$\text{Val}(A_6, b, I) = \text{norm}_{[0,1]}(\rho_{\vec{f}(I), \vec{f}(b)})$$

The value of A_6 is calculated as the Pearson's correlation coefficient of the vector-value property *DICinfo* of a blend b and an input I . Higher correlations in the DIC vectors of two transitions indicate higher resemblance; *norm* is a function that normalises the Pearson's coefficient from the interval $[-1, 1]$ to the interval $[0, 1]$.

For all the other types of arguments, Val is defined as:

$$\text{Val}(A_i, b, I) = \sum_{f_j \in \psi(A_i)} 1 - (f_j(I) - f_j(b))$$

Based on the above definitions, the *total association* value is the sum of the individual associations.

$$\text{assoc}(b) = \sum_{I_i \in \mathcal{I}} a(b, I_i)$$

where \mathcal{I} is the set of input spaces, containing in this specific case, I_1 and I_2 .

Symmetry, on the other hand, reflects the balance of properties that a blend inherits from both input spaces. A blend

has a high symmetry when it inherits an almost equal proportion of properties from both input spaces. Blends having higher symmetry are preferred to those with lower symmetry, since a high symmetry reflects a stronger hybridisation of structural characteristics. Hybridisation is an important principle to evaluate transition blends.

The blend symmetry is defined in terms of its 'asymmetry'. The asymmetry of a blend w.r.t. the inputs, denoted as $\text{asym}(b)$, is calculated as:

$$\left| \frac{a(b, I_1)^2 + a(b, I_1)a(b, I_2)}{a(b, I_1)^2 + a(b, I_2)} - \frac{a(b, I_2)^2 + a(b, I_1)a(b, I_2)}{a(b, I_2)^2 + a(b, I_1)} \right|$$

The value of $\text{asym}(b)$ is defined in $[0, 1]$, where 0 stands for a perfect symmetry (equal association with both inputs) and 1 stands for total asymmetry (association only with one input). Additionally, the non-absolute version of the above equation suggests the prevailing input, with a negative value indicating dominating association of the blend with the first input and a positive value contrarily.

The total rate of a blend is computed by taking the input association and asymmetry values into account.

$$\text{rate}(b) = \frac{\text{assoc}(b)(1 - \text{asym}(b))}{\text{assoc}(b) + (1 - \text{asym}(b))}$$

The above expression promotes pairs of association and symmetry that are both high, while a simple sum would allow a low value of the one to be covered by the other.

Finally, a decision making criterion to compare any pair of blends $b_1, b_2 \in \mathcal{B}$ can be defined as follows.

Definition 3 (Decision criterion). *A blend b_1 is preferred to a blend b_2 if and only if $\text{rate}(b_1) \geq \text{rate}(b_2)$.*

It is worthy to notice that the above criterion guarantees that any pair of blends is comparable, and, consequently, it allows to decide which blends are the best ones. This is an important property for blend evaluation and, generally, for approaches to argumentation-based decision making (Amgoud and Prade, 2009; Bonet and Geffner, 1996).

System evaluation

Referring to Figure 1, via the interface α , the user has access to modules γ , and β which can be used to specify *concepts* that will inform the resulting product, namely, the input idioms and arguments that impose constraints on the generated blend. These are translated by the system into process-friendly formats. Module ϵ embodies the (process-level) concept of a system that make use of the supplied idioms and the blending properties to generate *example* transition matrices, ζ . In the current version of the system, these transitions are evaluated by the user (music expert) in step ι using sophisticated harmonic knowledge that reflects an historically established musical *aesthetic*. The user can then return to the GUI α , and adjust the settings of γ and β to regenerate the transitions.

This is illustrated in Figure 5 in box **P1**, using the diagrammatic extension to the FACE model by Colton et al. (2014). Here, capital letters F , A , C , or E are creative

acts that generate a framing, aesthetic, concept, or example, respectively. Administrative acts S and T denote selection and translation. Lower-case letters denote the generated artefact in each case (e.g., the concept c corresponding to the concept-creation act C). Subscripts p , g , or m indicate whether the act takes place at the process, ground, or meta level. Inside each box, stacks show the dependence in development epochs, and arrows show run-time message passing. Acts taken by the programmer or user are decorated with a bar, whereas acts taken by the system itself receive no extra decoration.

In the current version of the system, apart from the *programmer's* creative acts specifying the modules and their interconnections, and the algorithm \overline{C}_p^ϵ that turns inputs into blends, the user, who is assumed to be a music expert, must intervene in the system in two places.

First, the *user* defines system settings \overline{C}_g^γ , \overline{C}_g^β that correspond to the selection of input idioms and of arguments respectively. Second, after the run completes, he or she evaluates the system output via \overline{A}_g^ϵ .

The *system's* primary responsibilities take place through the creative acts E_g^ϵ , which generate blends, and $S[a_g^{\beta,\gamma}](e^{\epsilon*})$, in which the aesthetic $A_g^{\beta,\gamma}$ (a unified label for assoc and asym, which are defined anew in each run, based on a fixed translation of the user's arguments, as specified in the previous section) is applied to rate the possible blends, and select to a final extended transition matrix.

Therefore, the key idea behind what has been implemented so far is an 'automated ranking/evaluation' step that guides the selection of blends, $S[a_g^{\beta,\gamma}](e^{\epsilon*})$ according to the arguments defined by the user. The development of the programmatic components that operationalise this process has relied on both computer science and musicological insights. This approach has been characterised as meaningful per se through informal feedback provided by musical experts – but is perhaps especially valuable because it constitutes a prototype for more involved automated evaluation of computer-generated harmonic spaces.

Indeed, the next step towards the development of a more autonomously creative system using the same architecture is fairly clear: future work would need to 'close the loop' computationally, connecting the evaluation of generated transition matrices with the parameter-setting (i.e., argumentation) stage, and making this run autonomously to refine the system's behavior. This as-yet hypothetical situation is illustrated in the box **P2**.

Here, the programmer has translated some of the user-specified aesthetics into code $\overline{T}[A_g^\epsilon]$, and invented a meta-level concept \overline{C}_m^α defining a system component that can automatically apply these aesthetics to the generated transition matrices e_g^ζ as in order to automatically generate new system settings C_g^γ , C_g^β .

Conclusion, Discussion and Future Work

In this paper, we described a methodology for harmonic blending and we proposed a creative system that assists musicologists with the evaluation and enhancement of harmonic innovation. We defined some harmonic features of




Figure 5: The current implementation **P1** prototypes automated evaluation of blends according to user's arguments; this points to a proposed future implementation **P2** with further automation.

chord transitions utilised for evaluating blends of transitions, leading to the invention of novel harmonic spaces. The system allows a musicologist to specify arguments over these features that are taken into account in the generation of new harmonic spaces. The music expert can then assess whether the new harmonic idiom makes musicological sense and re-adjust the arguments to explore alternative blends that can potentially produce better harmonic spaces.

The main advantage of the current system is the agile interaction through which the user can express desirable properties over the transition blends and their argument-based evaluation in order to produce musically meaningful results. The added value of argumentation is the ranking/evaluation of blended transition—obtained by conceptual blending of two input transition belonging to two musical idioms—by answering questions which abstract several properties of chord transitions. On the other hand, the evaluation of the creative output of the system, i.e., an extended harmonic space that includes blended transitions, is carried out by the user via an introspective argumentative dialogue.

In a future work we intent to use the argumentation-based process for evaluating the blended harmonisations of user defined melodies, i.e., actual music output. Additionally, mapping the properties of the blended idiom or, at a latter stage of a harmonised melody, back to the parameter-setting stage opens an interesting direction for future research and further improvements of the system. The added value of argumentation can be stressed, for instance, by letting the system suggest possible refinements of the initial user arguments, progressively converting part of the introspective user evaluation into a more explicit format. For example, a future version of the system would be based on identifying harmonic features of the input spaces that automatically suggest an 'optimal' set of initial arguments. The current version of the system is an already-usable prototype on the way towards the development of a more autonomous creative system.

Acknowledgments

This work is partially supported by the COINVENT project (FET-Open grant number: 611553).

References

- Amgoud, L., and Prade, H. 2009. Using arguments for making and explaining decisions. *Artificial Intelligence* 173(3-4):413–436.

- Antovic, M. 2009. Musical Metaphors in Serbian and Romani Children: An Empirical Study. *Metaphor and Symbol* 24(3):184–202.
- Antovic, M. 2011. Musical metaphor revisited: Primitives, universals and conceptual blending. *Universals and Conceptual Blending (February 17, 2011)*.
- Bench-Capon, T. J. M., and Dunne, P. E. 2007. Argumentation in artificial intelligence. *Artificial Intelligence* 171(10-15):619–641.
- Bonet, B., and Geffner, H. 1996. Arguing for decisions: A qualitative model of decision making. In *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence (UAI'96)*, 98–105.
- Cambouropoulos, E.; Kaliakatsos-Papakostas, M.; and Tsougras, C. 2014. An idiom-independent representation of chords for computational music analysis and generation. In *Proceeding of the joint 11th Sound and Music Computing Conference (SMC) and 40th International Computer Music Conference (ICMC)*, ICMC–SMC 2014.
- Cambouropoulos, E.; Katsiavalos, A.; and Tsougras, C. 2013. Idiom-independent harmonic pattern recognition based on a novel chord transition representation. In *In Proceedings of the 3rd International Workshop on Folk Music Analysis (FMA2013)*, FMA 2013.
- Cambouropoulos, E. 2012. A Directional Interval Class Representation of Chord Transitions. In *Proc. of the 12th International Conference for Music Perception and Cognition, & 8th Conference of the European Society for the Cognitive Sciences of Music*, ICMPC-ESCOM 2012.
- Colton, S.; Pease, A.; Corneli, J.; Cook, M.; and Llano, T. 2014. Assessing progress in building autonomously creative systems. In Ventura, D.; Colton, S.; Lavrac, N.; and Cook, M., eds., *Proceedings of the Fifth International Conference on Computational Creativity*.
- Confalonieri, R.; Schorlemmer, M.; Plaza, E.; Eppe, M.; Kutz, O.; and Peñaloza, R. 2015a. Upward Refinement for Conceptual Blending in Description Logic – An ASP-based Approach and Case Study in EL⁺⁺. In *International Workshop on Ontologies and Logic Programming for Query Answering, International Joint Conference on Artificial Intelligence (IJCAI) 2015*.
- Confalonieri, R.; Corneli, J.; Pease, A.; Plaza, E.; and Schorlemmer, M. 2015b. Using Argumentation to Evaluate Concept Blends in Combinatorial Creativity. In *Proceedings of the 6th International Conference on Computational Creativity, ICCCI5*, 174–181.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2):321 – 357.
- Eppe, M.; Confalonieri, R.; Maclean, E.; Kaliakatsos-Papakostas, M.; Cambouropoulos, E.; Schorlemmer, M.; Codescu, M.; and Kühnberger, K.-U. 2015a. Computational invention of cadences and chord progressions by conceptual chord-blending. In *International Joint Conference on Artificial Intelligence (IJCAI) 2015*.
- Eppe, M.; Maclean, E.; Confalonieri, R.; Kutz, O.; Schorlemmer, M.; and Plaza, E. 2015b. ASP, Amalgamation, and the Conceptual Blending Workflow. In Calimeri, F.; Ianni, G.; and Truszczynski, M., eds., *Logic Programming and Nonmonotonic Reasoning*, volume 9345 of LNCS. Springer International Publishing. 309–316.
- Fauconnier, G., and Turner, M. 2002. *The Way We Think: Conceptual Blending And The Mind's Hidden Complexities*. Basic Books.
- Goguen, J., and Harrell, D. F. 2010. Style: A Computational and Conceptual Blending-Based Approach. In Argamon, S., and Dubnov, S., eds., *The Structure of Style: Algorithmic Approaches to Understanding Manner and Meaning*. Berlin: Springer. 147–170.
- Goguen, J. 2006. Mathematical Models of Cognitive Space and Time. In Andler, D.; Ogawa, Y.; Okada, M.; and Watanabe, S., eds., *Reasoning and Cognition*, volume 2 of *Interdisciplinary Conference Series on Reasoning Studies*. Keio University Press.
- Kaliakatsos-Papakostas, M., and Cambouropoulos, E. 2014. Probabilistic harmonisation with fixed intermediate chord constraints. In *Proceeding of the joint 11th Sound and Music Computing Conference (SMC) and 40th International Computer Music Conference (ICMC)*, ICMC–SMC 2014.
- Kaliakatsos-Papakostas, M.; Zacharakis, A.; Tsougras, C.; and Cambouropoulos, E. 2015. Evaluating the General Chord Type representation in tonal music and organising GCT chord labels in functional chord categories. In *Proceedings of the 4th International Conference on Music Information Retrieval (ISMIR 2015)*.
- Ontañón, S., and Plaza, E. 2010. Amalgams: A formal approach for combining multiple case solutions. In *Proc. of the Int. Conf. on Case Base Reasoning*, volume 6176 of *Lecture Notes in Computer Science*, 257–271. Springer.
- Ritchie, G. D. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17:67–99.
- Schorlemmer, M.; Smail, A.; Kühnberger, K.-U.; Kutz, O.; Colton, S.; Cambouropoulos, E.; and Pease, A. 2014. Coinvent: Towards a computational concept invention theory. In *5th Int. Conf. on Computational Creativity*.
- Simon, I.; Morris, D.; and Basu, S. 2008. Mysong: Automatic accompaniment generation for vocal melodies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08*, 725–734. New York, NY, USA: ACM.
- Zacharakis, A.; Kaliakatsos-Papakostas, M.; and Cambouropoulos, E. 2015. Conceptual blending in music cadences: A formal model and subjective evaluation. In *Proceedings of the 4th International Conference on Music Information Retrieval (ISMIR 2015)*.
- Zbikowski, L. M. 2002. *Conceptualizing Music: Cognitive Structure, Theory, and Analysis*. Oxford University Press.