Conceptual blending of harmonic spaces for creative melodic harmonisation

Maximos Kaliakatsos-Papakostas[a], Marcelo Queiroz[b], Costas Tsougras[a], Emilios Cambouropoulos[a]

[a]School of Music Studies, Aristotle University of Thessaloniki
[b]Computer Science Department, University of São Paulo
{maxk,tsougras,emilios}@mus.auth.gr
mqz@ime.usp.br

Abstract

In computational creativity, new concepts can be invented through conceptual blending of two independent conceptual spaces. In music, conceptual blending has been primarily used for analysing relations between musical and extra-musical elements in composed music rather than generating new music. This paper presents a probabilistic melodic harmonisation assistant that employs conceptual blending to combine learned, potentially diverse, harmonic idioms and generate new harmonic spaces that can be used to harmonise melodies given by the user. The key feature of this system is the application of creative conceptual blending to the most common chord transitions (pairs of consecutive chords) of two initial harmonic idioms. The proposed methodology integrates newly created blended chords and transitions in a compound probabilistic harmonic space, that preserves combined characteristics from both initial idioms along with those new chords and transitions within a unified setting. This methodology enables various interesting music applications, ranging from problem solving, e.g. harmonising melodies that include key transpositions, to generative harmonic exploration, e.g. combining major-minor harmonic progressions or more extreme idiosyncratic harmonies.

1 Introduction

New concepts may be invented by traversing previously unexplored regions of a given conceptual space (exploratory creativity), transforming established concepts (transformational creativity), or by making associations between diverse conceptual spaces (combinational creativity); Boden maintains that the latter, i.e., combinational creativity, has proved to be the hardest to describe formally (Boden, 2009).

Conceptual blending is a cognitive theory developed by Fauconnier and Turner (Fauconnier and Turner, 2003) whereby elements from diverse, but structurally-related, mental spaces are combined, giving rise to new conceptual spaces: such spaces often possess new powerful interpretative properties allowing better understanding of known concepts or the emergence of altogether novel concepts. Conceptual blending is a process that allows the construction of meaning by correlating elements and structures of diverse conceptual spaces. It relates to Boden's notion of combinational creativity (Boden, 2009). A generative computational framework that incorporates the conceptual blending theory in its core model has been developed in the context of the Concept Invention Theory (COINVENT) project[1] (Schorlemmer et al., 2014).

With regards to music, conceptual blending has been predominantly theorised as the cross-domain integration of musical and extra-musical domains such as text or image (e.g. Tsougras and Stefanou (2015); Zbikowski (2002, 2008); Cook (2001); Moore (2013)), and primarily discussed from a musico-analytical perspective focusing on structural and semantic integration. Blending as a phenomenon involving 'intra-musical' elements (Spitzer (2004), Antovic (2011)) has received less attention. In principle, one of the main differences of blending theory from the theory of Conceptual Metaphor (CMT) is that it may involve mappings between incongruous spaces within a domain (e.g. conflicting tonalities in a musical composition). In this case, 'intra-musical' conceptual blending is often conflated in music with the notion of structural blending (Goguen and Harrell, 2010), and Fauconnier and Turner's theory is primarily applied to the integration of different or

[1]http://www.coinvent-project.eu/en/home.html

conflicting structural elements, such as melody notes of different pieces according to their time-span reduction (Hirata et al., 2014) chords, harmonic spaces, or even melodic-harmonic material from different idioms (e.g. Kaliakatsos-Papakostas et al. (2014); Ox (2014)). A more extended discussion and critical examination of conceptual blending processes in music is presented in (Stefanou and Cambouropoulos, 2015).

Different musical styles/idioms establish independent harmonic spaces that involve a network of inter-related constituent concepts such as chord, root, scale hierarchy, tonality, harmonic rhythm, harmonic progression, voice-leading, implied harmony, reduction, prolongation, and so on. Conceptual blending is facilitated when a rich background (Schorlemmer et al., 2014) of concepts is available and when these concepts are structured in such ways that creative mappings are supported. Thereby, the existence of a rich background that includes formal descriptions of diverse harmonic elements is required; the combination of concepts from different idioms injects novelty and creativity to the melodic harmonisation process.

Harmony has been modelled either with mathematical/geometrical models that provide extra-musical perspectives (Tymoczko, 2006; Callender et al., 2008), or, in the case of tonal and jazz harmony, effective models have been proposed that utilise grammar-related structures (Rohrmeier, 2011; Koops et al., 2013; Granroth-Wilding and Steedman, 2014). The concepts in the aforementioned mathematical/geometrical spaces can be used in a theory-driven approach to create new combinations of harmonic transformations (Callender et al., 2008), leading to new combined harmonic possibilities. However, for the data-driven purposes of blending different musical styles learned from data, it is necessary to utilise harmonic representations that are substantially different from the aforementioned ones. Toward this end, an idiom-independent representation of harmonic concepts, such as the General Chord Type representation, or GCT (Cambouropoulos et al., 2014)[2], is more adequate. This representation has already been used as the basis of a modular hierarchical representation of harmonic structure (Kaliakatsos-Papakostas et al., 2016b) that allows 'meaningful' blends at various hierarchic levels of harmony for practically any musical idiom. Knowledge extracted from a large dataset of more than 400 harmonically annotated pieces[3] (manually produced harmonic reductions) from various diverse musical idioms (from medieval to 20th century styles) comprise the rich background required for interesting and creative blends. More specifically, from a set of harmonic reductions for a given idiom (e.g. Bach chorales, tango songs, jazz standards, etc.) the following structural characteristics are learned/extracted: chord types, chord transitions (probabilistic distributions), cadences (i.e. chord transitions on designated phrase endings at different hierarchic levels), and voice-leading (i.e., bass line motion in relation to melody, bass-melody distance, chord inversion). Such features from diverse idioms may be combined giving rise to new harmonic blended styles; for instance, tonal cadences may be assigned to phrase endings and modal chord transitions may be employed for filling in the rest of the phrase chords – examples may be found in Cambouropoulos et al. (2015).

This paper focuses on the following questions: Can chord transitions per se be blended? Can two different chord transitions (e.g. cadences) from different idioms be combined to give rise to novel transitions that do not appear in any of the input harmonic spaces? Additionally, can whole chord transition matrices from different harmonic styles be amalgamated so as to generate new chord transition spaces? The developed methodology presented in this paper[4] proposes a new approach to creative musical systems that incorporates *learning* from data and *blending* of learned elements to create new harmonic spaces. In contrast to systems that are able to mimic specific musical styles by learning from data (e.g. see Raphael and Stoddard (2004); Dixon et al. (2010); Whorley et al. (2013)), or by encoding musical knowledge of specific musical idioms by expert-designed rules (e.g.

[2]For instance, in a C major scale, the GCT representation of a C major chord is [0, 0 4 7], meaning that pitch class 0 is the root and chord notes are obtained as $0 + 0$, $0 + 4$ and $0 + 7$; a G7 chord is [7, 0 4 7 10], revealing the underlying structure [0 4 7] of any major chord simply shifted to the root 7, whereas a B full diminished chord is easily seen to correspond to the GCT representation [11, 0 3 6 9]. The reader interested in the use of the GCT representation in idiom-independent music contexts is referred to (Cambouropoulos et al., 2014) for more details.

[3]`http://ccm.web.auth.gr/datasetdescription.html`

[4]This methodology is incorporated in the CHAMELEON (Creative Harmonisation of Melodies via Learning and Blending of Ontologies) melodic harmonisation assistant. The interested reader is referred to the following webpage: `http://ccm.web.auth.gr/chameleonmain.html`.

see Ebcioglu (1988); Phon-Amnuaisuk et al. (2006)), the presented melodic harmonisation assistant is able to 'extrapolate' from the learned structures by blending elements learned from data. Thereby, new structures are created that preserve the most meaningful structural relations of the blended styles, generating new spaces that incorporate new characteristics. These structures incorporate information about the pitch classes of chord transitions. The core mechanism described in this paper is incorporated in a melodic harmonisation assistant that provides the user with harmonisations that come out from blended harmonic spaces with rudimentary voice leading, which requires human interventions (more or less severe) for becoming a 'correct' harmonisation, according to the voice leading rules of a specific musical style.

The core mechanism for blending harmonic styles is a methodology that blends isolated learned chord transitions, in order to create new harmonic spaces that can be used to harmonise melodies. Chord transition blending in the special case of cadence blending, has been explored in previous studies (Eppe et al., 2015a; Zacharakis et al., 2015). In these studies, two cadences (e.g. the tonal Perfect cadence and the modal Phrygian cadence) that share the same final tonic chord are blended giving rise to new chord progressions (e.g., the Tritone Substitution cadence that is commonly employed in jazz); the generated new cadences feature important characteristics from both of the input cadence spaces, namely ascending and descending leading notes to the tonic, preserving thus the closure effect of the resulting 'new' cadential formulae. In this paper, the cadence blending process (which is based on the COINVENT conceptual blending mechanism – see Section 2 for a brief description) is generalised to any two input chord transitions, allowing the creative blending of entire chord transition matrices from different idioms.

Let us attempt to illustrate the above chord transition blending processes by employing a simplistic harmonic blending example, whereby the blended spaces are merely different diatonic major tonalities. Suppose one has available (manually constructed or learned) a purely diatonic hidden Markov model on the C-major scale, with a first order chord (state) transition matrix and diatonic observed melodies. If a newly given C major melody contains a harmonically structural F♯ note, then the Markov model reaches a dead-end as it does not know of any diatonic C major chord that can harmonise this chromatic note. If two neighbouring tonalities, however, are blended, i.e. C major and G major, then the resulting composite transition matrix contains the D major chord that leads as the dominant to the tonic in G major or as secondary dominant to the dominant in C major (see Section 3 below). For a major tonality, borrowed chords from the relative or parallel minor keys and from neighbouring tonalities can be seen as single-scope blends (following the terminology of Fauconnier and Turner (2003)), i.e., blends in which one primary input space remains mostly intact and specific features are imported from the secondary space.

A more extreme blend would occur for instance between the C major and F♯ major tonalities represented by purely diatonic Markovian spaces. Since these spaces have no common diatonic chords, the two transition matrices for these tonalities do not 'overlap' at all, and there is no way to make the transition from one space to the other. In such a case, chord transition blending may be employed to try to find new potential chord transition candidates that may allow an 'acceptable' transition between the two spaces. Let us assume that only three basic chords for each space are available, namely the tonic, subdominant and dominant seventh major chords of each tonality; the transition matrices for these two 'toy' spaces do not communicate, since there is no path for transiting form chords of one space to the chords of the other. This is indicated by the zero probability values on top-right and bottom-left squares in Figure 1 (a) (where the colour of each box corresponds to a probability value and zero probability is indicated by white colour; darker colours indicate higher probability values). Can the proposed chord transition methodology 'invent' new transitions and potentially new chords that may connect the two spaces in a meaningful way?

The chord transition blending methodology is applied to all the chord transitions in the C major and F♯ major tables, i.e. each chord transition in the first matrix is blended with each chord transition in the second matrix producing a list of resulting blends. The resulting blends are ranked according to certain criteria that take into account the number of common features shared by the input chord transitions preserved in the blend. The features that are taken into account include common pitch classes in the first and/or second chords of the blend in relation to the two input transitions, common ascending and/or descending semitone movements in the transitions, and as-

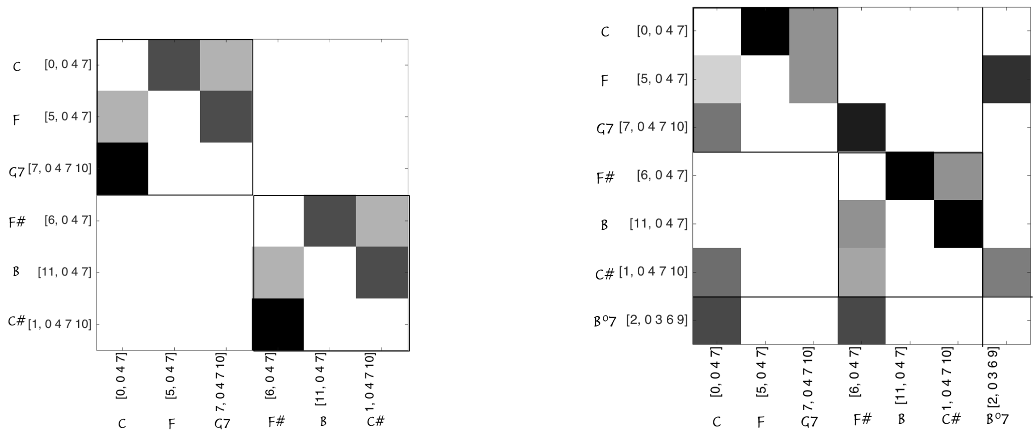

(a) No blends

(b) Topmost blends

Figure 1: Simple C major and F♯ major harmonic transition spaces with (b) no transition blends involved and (b) incorporating some of the topmost blending transitions.

cending and/or descending semitone movements to the root of the final chord of each transition (see detailed description in Section 2). For instance, the transition G7 → C and C♯7 → F♯ share the same pitch classes (pcs) 5 and 11 in the first chord, have similar ascending and descending semitone movements between the two chords and contain an ascending semitone movement to the root of the final chord.

By learning GCT and chord transition representations from datasets and generalising on properties representing chords and chord transitions, this methodology is able to produce novel harmonic structures: e.g. by generalising on the root property while keeping a given chord type, one easily obtains all 12 circular shifts of that chord type (many of which might not appear in the input dataset); and this is just one of the possible generalisations. Chord structure dictionaries can thus be learned from data and used to produce new blended chord transitions. Returning to the C major and F♯ major example, and assuming that we have available a palette of basic chord types, e.g. major, minor, major seventh, diminished and diminished seventh chords, a chord transition blend that ranks high is a transition in which the first chord is a diminished seventh (pcs: 2 5 8 11) and the second chord is either of C or F♯ (among other things the diminished seventh shares two common pcs with each of the first chords of the input transitions). The system is indeed able to 'invent' this diminished chord (pcs: 2 5 8 11), even if the initial harmonic spaces do not include it, and rate some blended transitions that incorporate this chord as good blends. Another good blend is one where the first chord is a major seventh chord a semitone above the tonic of each space (e.g. 1 5 8 11), which is a kind of tritone substitution transition; in this blended transition both chords found in the initial spaces before blending, but no transition between them existed in the original models. These newly invented transitions are illustrated in the new grey boxes added in the matrix of Figure 1 (b).

As seen in the above example, chord transition blending can be employed to create new transitions (and potentially new chords in these transitions) that preserve important features of the input transitions. When only the top ranking blends are preserved, then the system has introduced a way to connect the two input chord spaces. If more blends are selected then the composite transition matrix becomes more populated allowing more connections between the spaces. If the probabilities of the new 'invented' transitions are low, then the chord generation system creates chord sequences mostly within each of the constituent input spaces occasionally allowing passage from one to the other. If the probabilities of the new blended transitions are increased, then the whole space becomes unified and movement between most or all of the chords of both spaces is enabled. This latter strong blending between input spaces can generate new harmonic spaces that are radically different from the initial input spaces (e.g. blending two diatonic major tonalities in different keys may give rise to a composite blended space that features strong chromaticism reminiscent of music appearing centuries after diatonic tonality – see examples in Section 4).

The proposed blending paradigm seems to introduce an intelligent way to address the traditional problem of zero probability transitions in Markov models (Cleary and Teahan, 1995). Rather than assigning arbitrary non-zero 'escape' probability values (Chordia et al., 2010) or enforcing arc-consistency (Pachet et al., 2011) to allow a Markov process to cope with cases it has not seen in the training data, different transition matrices can be blended (or even a single matrix can be blended with itself) in order to introduce transitions that preserve qualities of the already existing transitions. At least for music, this seems to be a reasonable way to bypass the problem of sparse input data (e.g. learning transitions of pitch or chords or rhythmic values from a single piece rather than from a large homogeneous dataset).

In the sections below, the COINVENT blending core model will be first presented, in order to show how it is applied to chord transition blending. Then, the chord transition matrix blending methodology will be described. Afterwards, a number of potentially interesting examples illustrating harmonic blending in melodic harmonisation will be given, through the presentation and discussion of melodies harmonised in different idioms and blends between these idioms. These results present different cases where harmonic blending can be useful, either as a problem solving or as a creative tool. The new possibilities offered in automated melodic harmonisation by the presented system indicate the overall usefulness of the COINVENT framework for inventing new concepts through conceptual blending. Finally, a brief overview of methods developed to evaluate the proposed methodology is given; evaluation of the methodology is performed both under the scope of how the resulting harmonisations are perceived, and to what extent these results are capable of enhancing human creativity. Specifically, results obtained on the perceptual characteristics of the products of this methodology, presented in (Zacharakis et al., 2017), indicate that blending two harmonic spaces results in melodic harmonisations that are either perceived as belonging to a harmonic style lying in between the two combined spaces, or as belonging to a new yet intrinsically related harmonic style, fulfilling the intended purposes of blending.

2 A computational framework specialised for blending chord transitions

In computational creativity, conceptual blending has been modelled by Goguen (2006) as a generative mechanism, by describing input spaces as *algebraic specifications* and computing the blended space as their categorical *colimit*. A computational framework that extends Goguen's approach has been developed in the context of the COINVENT[5] (Concept Invention Theory) project (Schorlemmer et al., 2014). According to this framework, two *input spaces* are described as sets of features, properties and relations and after their *generic space* is computed, an *amalgamation* process (Eppe et al., 2015b; Confalonier et al., 2015) leads to the creation of several blends, which can be ranked in terms of value according to some criteria that relate to the knowledge domain of the modelled spaces.

In conceptual blending the properties of two input conceptual spaces are combined to create new spaces. The input spaces share some common structure along with differences. The intended goal of conceptual blending is to achieve a 'meaningful' combination of the non-common structural parts so that new structure emerges, giving novel properties to the generated blended space. An important aspect of the blended space is to preserve the *common* parts of the input spaces. The *generic space* is the conceptual space that keeps the common structure of the input spaces, guaranteeing that this structure also exists in the blended space, and generalises or abstracts over the parts of the input spaces that are distinct. In case-based reasoning the generic space is described as 'the most specific generalisation' (Ontañón and Plaza, 2012) of the input spaces, in the sense that the generic space is obtained as the first common element in a hierarchy of generalisations built on top of each input space, so that the inputs are generalised by the least possible 'amount' and only their common parts

[5]http://www.coinvent-project.eu

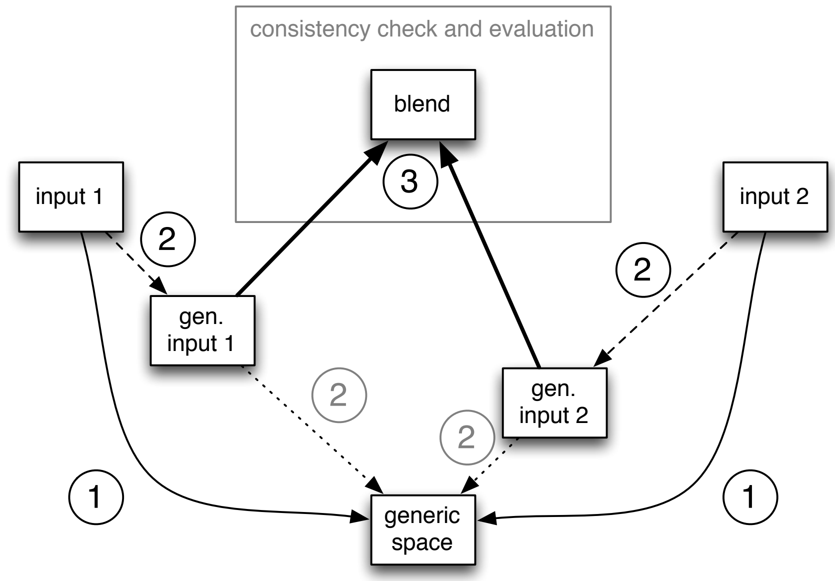

Figure 2: Conceptual blending based on amalgamation. The generic space is computed (1) and the input spaces are successively generalised (2), creating successively new potential blends (3). Some blends might be inconsistent or poorly evaluated according to ranking criteria or domain specific criteria.

remain in this generic space (i.e. generalising more would also remove parts that are common to both input spaces).

2.1 The COINVENT framework for conceptual blending

The COINVENT framework for generative conceptual blending is based on the notion of *amalgamation* and is illustrated in Figure 2. This process generates *amalgams* or *blends* of two input spaces, which are roughly new spaces that contain parts from the initial ones (Confalonier et al., 2015). For reasons of clarity and text consistency, in the following paragraphs the term 'blend' will be used for referring to the products of amalgamation. 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. Regarding blends, the terms 'consistent' and 'inconsistent' refer to whether all logical relations in the blend and the background knowledge are satisfied, i.e. there are no mutually cancelling contradictions. Figure 2 illustrates the amalgam-based COINVENT algorithmic model for conceptual blending.

Amalgam-based conceptual blending has been applied to invent chord cadences (Eppe et al., 2015a; Zacharakis et al., 2015). In this setting, cadences are considered as special cases of chord transitions – pairs of successive chords, occurring before a phrase/section boundary – that are described by means of properties such as the roots or types of the chords, or specific voice motions. 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 tonal perfect cadence with the renaissance Phrygian cadence (see Eppe et al. (2015a); Zacharakis et al. (2015)) — described by the transitions (Input 1) I_1: G7 $\rightarrow$ C and (Input 2) I_2: B♭m $\rightarrow$ C^5 respectively — their generic space would consist of any transitions that have a first chord with pitch class 5 (common to G7 and B♭m), a second chord with pitch class 0 and 7, where the first chord has a pitch class a semitone higher or lower than the second chord's root, possibly along with other properties generalised from the given description of the input transitions.

After a generic space is found, the amalgam-based process computes the blend of two input spaces by *unifying* their content. If the resulting blend is inconsistent, then it iteratively *generalises* (Ontañón and Plaza, 2012) 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 unify directly the transitions I_1: G7 → C and I_2: B♭m → C^5 would yield an inconsistent blend, since a transition cannot both include and *not* include an upward leading note to the second chord's tonic (which are features of I_1 and I_2 respectively, as discussed in more detail later). Therefore, the amalgam-based process generalises the clashing property with respect to one of the inputs (e.g., the property describing the absence of leading note in I_2 would be excluded) 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). In this specific cadence blending example, one novel blend that arises from the perfect and Phrygian cadences is the Tritone Substitution progression/cadence (that is commonly used in jazz). 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. For instance, the A♭ note in the tritone substitution invention example discussed in Eppe et al. (2015a) and Zacharakis et al. (2015) is imported through completion since the C♯7 chord is required to have a perfect fifth according to their cadence formalisations, where both input chord types have a perfect fifth.

The methodology for transition blending described in the paper at hand uses an equivalent to the aforementioned methodology that combines amalgamation and completion. Chords are represented using the General Chord Type (GCT) representation (Cambouropoulos et al., 2014). The proposed methodology is adjusted for the specific harmonic ontology (with the GCT representation), using a *dictionary* of chord types that are allowed in the emerging blends. This dictionary depends on the idioms that take part in the blending process and is learned from data, representing a part of the "background knowledge" that these idioms incorporate. Therefore, based on the assumption that only certain chord types are allowed, the search space of possible chords in blended transitions is not overwhelmingly large, thus for the specific task of transition blending the importance of the amalgamation process is reduced and can be omitted altogether. This modification is presented in detail in Section 2.2.

After several blends have been computed, an evaluation process ranks them according to criteria that reflect the importance of the properties that blends inherit from the input spaces. Blending quality is a necessary aspect of conceptual blending since it allows the identification of better blends among the many (potentially too many) possible ones[6]. In the general context of conceptual blending, several blending optimality principles have been proposed for rating and ranking blends (see e.g. Chapter 16 of Fauconnier and Turner (2003)), but a detailed description of optimality principles is unnecessary for the transition blending methodology here presented (the reader is referred to Goguen and Harrell (2010) for applications of several such principles in the *Alloy* algorithm). The proposed methodology for rating and ranking blends used in this paper is based on criteria concerning the salience of transition features within their idioms and is described in detail in Section 2.3.

2.2 Formal description and chord transition blending

A formal ontology of transitions is required for blending according to the COINVENT framework. A chord transition (a sequence of two chords) is described as a set of properties that involve each chord independently and the chord transition as a whole (relations between the two chords). In Kaliakatsos-Papakostas et al. (2016a), an argument-based system was presented that allowed music experts to define which transition properties should be considered, through observation of blending results obtained in various harmonic setups. Using the aforementioned argument-based system and after examination of several produced outcomes, a (non-conclusive) list of nine important properties was maintained:

1. *fromPCs*: the pitch classes included in the first chord,
2. *toPCs*: the pitch classes included in the second chord,
3. *DIChas0*: Boolean value indicating whether the Directed Interval Class (DIC) vector (Cambouropoulos, 2012; Cambouropoulos et al., 2013) of the transition contains a 0 (i.e. that both

[6]The amalgamation process produces many blends by following alternative generalisation paths.

chords have at least one common pitch class),

4. *DIChas1*: as above but for DIC value 1 (i.e., the transition contains at least one ascending semitone),
5. *DIChasMinus1*: as above but for DIC value −1 (i.e., the transition contains at least one descending semitone),
6. *ascSemNextRoot*: Boolean value indicating whether the first chord has a pitch class with ascending semitone relation to the pitch class of the second chord's root,
7. *descSemNextRoot*: as above but with descending semitone, and
8. *semNextRoot*: as above but with either ascending or descending semitone.
9. *5thRootRelation*: Boolean value indicating whether the first chord's root note is a fifth above the root of the second. Root notes of chords are computed with the General Chord Type (GCT) (Cambouropoulos et al., 2014) algorithm.

Table 1 demonstrates the property values (also referred to as features) for the three transitions (namely the perfect, Phrygian and tritone substitution cadences) of the example discussed above. In this example, the tritone substitution cadence has been produced as a result of blending between the perfect and the Phrygian cadential progressions, with a process that is described below.

Table 1: Blending the chord transitions of the minor-mode perfect cadence (GCT notation of Input 1: $[7, 0\ 4\ 7\ 10] \rightarrow [0, 0\ 3\ 7]$) and the Phrygian cadence (GCT notation of Input 2: $[10, 0\ 3\ 7] \rightarrow [0, 0\ 3\ 7]$). The common elements of both input spaces that are included in the generic space are depicted in boxes, while the other common elements are shown in circles. Many blends are produced by blending these cadences; the tritone substitution blend is shown in the last column of the table as an illustrative example.

Property name	Input 1 (Perfect)	Input 2 (Phrygian)	Possible blend
fromPCs	$\{7, 11, 2, \boxed{5}\}$	$\{10, 1, \boxed{5}\}$	$\{1, 5, 8, 11\}$
toPCs	$\{\boxed{0}, \boxed{3}, \boxed{7}\}$	$\{\boxed{0}, \boxed{3}, \boxed{7}\}$	$\{0, 3, 7\}$
DIChas1	1	0	1
DIChasMinus1	0	1	1
DIChas0	1	0	0
ascSemNextRoot	1	0	1
descSemNextRoot	0	1	1
semNextRoot	⓵	⓵	⓵
5thRootRelation	1	0	0

In the COINVENT framework for computational conceptual blending, the role of the generic space, which includes all common elements of the input spaces, is to reject possible blends that do not incorporate these common elements. After extensive experimentation during the development of the presented transition blending methodology, it became obvious that a richer representation of transitions (i.e. one that incorporates many properties) potentially led to stricter generic space demands (i.e. generic spaces with more fixed properties), allowing a smaller number of 'surprising' blends to be generated. The generic space requirements are necessary for discarding blends that do not capture the important common features from the input spaces. To this end, two types of properties are distinguished: the *necessary* and the *desired* properties of transition blending. *Necessary* properties are elements incorporated in the generic space, i.e. if a necessary property is common in both inputs, then blends that don't have it are rejected. *Desired* properties are properties that characterise the input spaces and are preferred to be part of a blend, but do not belong to the generic space (i.e. they are not necessarily included in every blend). Both necessary and desired properties play an important role in rating and ranking the blends as described in Section 2.3.

In the context of the current study, among the nine properties that describe transitions, only two that concern the pitch classes of the involved chords are considered as *necessary*, namely the *fromPCs* and *toPCs*. The example in Table 1 demonstrates the role of the necessary and desired properties in transition blending. Therein, boxed items indicate the common elements in the input transitions regarded as necessary properties. For instance, all the pitch classes of the second chord as well as pitch class 5 in the first chord are present in both inputs and, therefore, are also included in all possible blends. On the other hand, the *desired* property *semNextRoot* is common in both inputs (indicated by circled numbers in Table 1); blends that do not include this property are allowed, but their rating will probably be low, depending on the salience value of this property in the inputs, to be discussed later.

So far, the discussion revolved around describing transitions with necessary and desired properties, but how are blends actually created? According to the amalgamation process, features from the input spaces should be successively generalised up to the point where no contradicting material is included (see Section 2.1 and Figure 2). This process is computationally expensive, since there are multiple generalisation paths that can be followed. Furthermore, additional musical criteria are required to check whether the generated blends are transitions that include 'acceptable' chord types – it is possible for the algorithm to generate note clusters or trivial single-note chords that haphazardly satisfy the generic space requirements and achieve high rating value.

Employing a dictionary of acceptable chord types automatically learned from data, it is possibly to omit the amalgamation process in order to explore more efficiently the space of possible blends. By assuming that the dictionary of chord types, denoted by $\mathcal{T}$, consists of N chord types, then all the possible chords that have to be examined are $12\ N$ – every chord type with every pitch class as a root note (i.e. all transpositions). Thereby, the 'universe' of all transitions between acceptable chords are $144\ N^2$. The number of acceptable types (N) is not overwhelmingly large for most musical idioms. For instance, by considering major and minor chords along with their sevenths, plus the half and full diminished chords (6 types in total), 5184 possible transitions can be generated. Therefore, producing valid blends is not a matter of constructing the proper chord types, but finding the already existing types that when assigned a proper root, satisfy the necessary and desired attributes. The algorithm for constructing the list of all valid blends is described in Figure 3. For computing all possible valid blends between two input transitions, first the generic space of the inputs is computed (`getGenericSpace(`I_1`,`I_2`)` in line 2, as described in section 2.1); then all acceptable chords are constructed (in line 4, as explained in section 1); all transitions between all possible chords are formed (in line 10 by considering all pairs of acceptable chords); and finally, each transition is examined (`satisfies(`tr`, `g`)` in line 13) in terms of its compatibility with the generic space produced by the input transitions, in order to verify all generalised properties defining the generic space so that this transition can be inserted in the list of acceptable blends. If a transition does not satisfy the generic space requirements, i.e. if it fails to verify a property common to both inputs, it is rejected as a possible blend. Section 2.3 analyses the process of rating and ranking all transitions in the list of acceptable blends.

As an example of running the algorithm in Figure 3, let us consider that the input transitions to the algorithm are the following:

$I_1 = [7,\ 0\ 4\ 7\ 10] \rightarrow [0,\ 0\ 4\ 7]$

$I_2 = [1,\ 0\ 4\ 7\ 10] \rightarrow [6,\ 0\ 4\ 7]$

and the dictionary of acceptable chord types is:

$\mathcal{T} = \{[0\ 4\ 7],\ [0\ 4\ 7\ 10], [0\ 3\ 6\ 9]\}$.

Running the function in line 2, we get the following properties for the generic space of those inputs:

{ *fromPCs*= {**5**, **11**}, *toPCs*= **ANY** } ← `getGenericSpace(`I_1`,`I_2`)`.

Therefore, it is necessary for every possible transition blend to incorporate the pitch classes 5 and 11 in the first chord, while it may incorporate any pitch class in the second chord (since there are no pitch class overlaps in the second chords of the inputs). Running the loop in line 4 creates the list of chords with all 12 possible pitch classes as roots and types that are included in $\mathcal{T}$. The loop in line 10 creates all the transitions between them. In order to examine whether two example

Algorithm 1 Computation of all possible blends

Require: (**i**) two input transition, I_1 and I_2, (**ii**) a dictionary of all acceptable chord types $\mathcal{T}$
Ensure: List of all possible blends ($\mathcal{B}$) of I_1 and I_2

$\mathcal{B} \leftarrow \emptyset$ {% initialise and empty set of blends}
$g \leftarrow$ `getGenericSpace(`I_1`,`I_2`)` {% get the generic space of inputs}
$\mathcal{C} \leftarrow \emptyset$ {% initialise the set of all possible acceptable chords}
{% make the set of all possible acceptable chords}
for $t \in \mathcal{T}$ **do**
  for $r \in \{0, 1, \ldots, 11\}$ **do**
    c = `makeChordWithRootAndType(`r`,`t`)` {% create a new chord with the currently selected root and type}
    $\mathcal{C}$ = `append(`$\mathcal{C}$`,`c`)` {% put this chord in the current list of chords}
  end for
end for
{% for all chord pairs}
for $c_1 \in \mathcal{C}$ **do**
  for $c_2 \in \mathcal{C}$ **do**
    tr = `formTransition(`c_1`,`c_2`)` {% form the transition from c_1 to c_2}
    {% check if transition satisfies generic space}
    if `satisfies(`tr`, `g`)` **then**
      $\mathcal{B}$ = `append(`$\mathcal{B}$`,`tr`)` {% if transition satisfies generic space, put it in the blends list}
    end if
  end for
end for

Figure 3: Algorithm for obtaining all possible transition blends of two input transitions, given a dictionary of learned acceptable chord types.

transitions/possible blends (tr_1 and tr_2) among the ones in the constructed list satisfy the generic space space requirements in line 13, let us consider the following two generated transitions as an example:

tr_1**:** [2, 0 3 6 9] → [0, 0 4 7]

tr_2**:** [4, 0 4 7] → [0, 0 4 7].

The property values of these transitions that are examined for generic space inclusion (*fromPCs* and *toPCs* properties) are the following:

tr_1**:** { *fromPCs*= {2, **5**, 8, **11**}, *toPCs*= {**0**, **4**, **7**} }

tr_2**:** { *fromPCs*= {4, 8, **11**}, *toPCs*= {**0**, **4**, **7**} }.

We can see that the tr_1 includes all the necessary generic space elements (shown in bold) while tr_2 fails to include the necessary pitch class 5 in its first chord. Therefore, tr_1 is included in the final list of possible blends ($\mathcal{B}$), while the tr_2 is discarded.

2.3 Rating a blend

The algorithm described in Figure 3 produces a list, $\mathcal{B}$, that includes all possible acceptable transitions that are potential blends of two given input transitions (I_1 and I_2). All blends in $\mathcal{B}$ need to be rated and ranked so that *meaningful* blends are distinguished and considered with higher priority for the next steps described in Section 3. When blending two transitions taken from two different harmonic spaces, the most *meaningful* blends would expectedly include a combination of all the salient features that the input transitions encompass. The salience of a feature of a transition, however, depends on the idiom that this transition belongs to. For a set of transitions in a certain harmonic context, the more rare or characteristic a feature is, the more salient/prominent it is considered. For instance, in C major the note transition B→C (11 → 0) appears in and characterises fewer chord transitions (namely G→C and Bdim→C), than e.g. note transition G→A (7 → 9) which appears in more transitions (e.g. G→Am, C→Am, C→F, C→Dm, G→F).

To compute the salience of a feature in a transition taken from an idiom, the above mentioned 'uniqueness' of this feature needs to be quantified. To this end, let us consider the set of all transitions in an idiom, denoted by $T_{\mathcal{I}}$, where $\mathcal{I}$ is the set of indexes of all transitions in the examined idiom. Also let T_i, $i \in \mathcal{I}$ be a transition from the examined idiom. Each transition property is considered as a function of a transition, $F_p(T_i) = v_p$, returning the value of this property in a specific transition – denoted by v_p. For instance, if T_i is the perfect cadence transition (G7 → C) and $F_{\text{ascSemNextRoot}}$ is the binary function returning the *ascSemNextRoot* property value (0 or 1 for not having or having an ascending semitone to next root, respectively), then the value of this property in the perfect cadence transition is obtained by $F_{\text{ascSemNextRoot}}(T_i) = 1$. We define the set of all transitions having a property p with a value v_p as

$$P_{p=v_p}(T_{\mathcal{I}}) = \{T_i, i \in \mathcal{I}; F_p(T_i) = v_p\},$$

while the cardinality (number of elements) of this set is denoted as $\mathcal{C}(P_{p=v_p}(T_{\mathcal{I}}))$. The *salience* of a property value v_p in a transition is therefore *inversely proportional* to the number of all transitions in the idiom that also include this property value. Hence, the *salience*, denoted by $S_{p=v_p}(T_i)$, of a property value v_p of a transition T_i is computed as

$$S_{p=v_p}(T_i) = \frac{1}{\mathcal{C}(P_{p=v_p}(T_{\mathcal{I}}))}.$$

This salience is only defined for values v_p appearing in some transition in the idiom, from which it is immediate that the denominator above does not vanish and the feature is well-defined.

An example of applying this methodology for computing saliences is given in Table 2. The considered training idiom in this example is a set of Bach chorales in major mode, after performing

GCT-based grouping (Kaliakatsos-Papakostas et al., 2015) of chords. Specifically, only the 10 most frequently used transitions of this idiom are considered, which represent each idiom as analysed in Section 3. The transitions incorporated in the example are [7, 0 4 7] → [0, 0 4 7] and [11, 0 3 6] → [0, 0 4 7] and the examined saliences concern the values of the *fromPCs* property. Both transitions include pitch class 11 as a *fromPCs* property value, but since they are the only transitions among the 10 ones representing the idiom that have it, the total salience of this feature is equally distributed among these two transitions (the 11 value of their *fromPCs* property has salience 0.5). Contrarily, the other *fromPCs* values are given smaller salience values, since they are also found in other transitions. For instance, *fromPCs* value 7 is assigned a value of 0.20 since it is shared in the first chords of five transitions in total (i.e. $\mathcal{C}(P_{p=v_p}(T_{\mathcal{I}})) = 5$ in the previous equation); except from [7, 0 4 7] → [0, 0 4 7], it is also found in 4 additional transitions that include [0, 0 4 7] as a first chord (leading to [7, 0 4 7], [5, 0 4 7], [2, 0 3 7] and [9, 0 3 7]). Similarly *fromPCs* value 2 is shared by three and value 5 is shared by four transitions.

Table 2: Example of saliences in the respective *fromPCs* property values of two transitions in a set of major mode Bach chorales: [7, 0 4 7] → [0, 0 4 7] and [11, 0 3 6] → [0, 0 4 7]. Since pitch class 11 appears as a member of the first chord only in these two transitions, the total salience of pitch class 11 in the entire idiom is equally distributed among them.

Idiom trained on a set of major-mode Bach chorales		
example transition:	[7, 0 4 7] → [0, 0 4 7]	[11, 0 3 6] → [0, 0 4 7]
fromPCs property values:	{7, **11**, 2}	{**11**, 2, 5}
respective saliences within idiom:	{0.20, **0.5**, 0.33}	{**0.5**, 0.33, 0.25}

A rating value is attributed to each blend in $\mathcal{B}$ for ranking them. The rating value of a blend in $\mathcal{B}$ is computed by summing all the saliences of features that this blend inherits from the input spaces. This sum is related to the harmonic mean of the cardinalities $\mathcal{C}(P_{p=v_p}(T_{\mathcal{I}}))$ above; in fact it is precisely its reciprocal times the number of common features. Therefore, blends that incorporate a larger total of salience values inherited from the inputs are ranked as better blends, while blends that inherit either fewer or less-salient features are ranked as worse blends.

3 Blending harmonic spaces via chord transition blending

The chord transition blending methodology described in Section 2 is integrated into the melodic harmonisation assistant presented in Kaliakatsos-Papakostas et al. (2016b). This assistant combines several probabilistic modules that learn musical structures from data, including chord transitions and cadences; chords are encoded using the GCT algorithm (Cambouropoulos et al., 2014), transitions are learned and composed with the constraint hidden Markov methodology (cHMM) (Kaliakatsos-Papakostas and Cambouropoulos, 2014), statistical models define the bass line voice leading (Makris et al., 2015) and a module fills the inner chord voices. The problem of assigning voice leading in chords with arbitrarily many voices, in an idiom-independent manner, has not yet been studied thoroughly in the literature, and is a problem that incorporates extensive combinatorial complexity. The method proposed in Kaliakatsos-Papakostas et al. (2016b) and utilised in this paper focuses mainly on fixing the bass voice given the melody, leaving relatively few possible combinations for the inner voices, given that these voices need to instantiate given pitch classes. Therefore, the probabilistic methodological attempt to solve the difficult and relatively unexplored problem of assigning voice leading on chords with arbitrary number of voices potentially leads to harmonisations that at some points might not adhere to strict voice leading rules; however, this method has the advantage of idiom-independency. It should be remarked that the focus of this paper is to explore the possibilities opened by using blending on the more abstract level of chord transitions, when chords are viewed in the form of pitch classes (through the General Chord Type), and so voice leading issues have been treated as a side aspect and dealt with as means to better illustrate the applications of transition blending.

This study attempts to employ chord transition blending in the context of the cHMM algorithm, with a view to combining *creatively* the independent chord transition matrices of two different harmonic idioms; the term ‘creatively’ here is used to indicate that the result of this process is a *novel* and *consistent* (in terms of musical criteria incorporated in chord transitions and in the learned harmonic spaces) composite harmonic space. To this end, GCT chord Markov transition tables learned from two *initial* idiom datasets are employed and most common chord transitions are indicated. Afterwards the transition blending methodology is applied on pairs of the most common transitions across the initial idioms, producing new blended transitions that connect and extend the possibilities of the initial idioms, generating a *compound* idiom that preserves some characteristics (in terms of transition probabilities) of the input idioms. Before transition blending is applied, a methodology for identifying common or similar chords of the initial idioms is employed; this enables potential connections between the two transition tables, making common-sense musical connections between the initial idioms (see Section 3.2).

The cHMM methodology used in Kaliakatsos-Papakostas et al. (2016b) incorporates a first order Markov model, indicating the probabilities of transitions from one chord to all possible next ones, and is here extended to include probabilities for chord-to-pitch classes relations and a method for dealing with intermediate chord constraints. Focusing only on transition probabilities, a convenient way to represent a first order Markov model uses transition matrices, including one row and one column for each chord in the examined idiom. The value in the i-th row and j-th column corresponds to the probability of the i-th chord transitioning to the j-th chord – the probabilities of each row sum to one. Figure 4 illustrates a colour-based graphic representation of the transition matrix learned from a collection of Bach Chorales in major mode (darker cells indicate higher probabilities). The displayed chords are actually GCT chord groups obtained by the method described in Kaliakatsos-Papakostas et al. (2015), while transitions between chords that pertain to the same GCT chord group are disregarded.

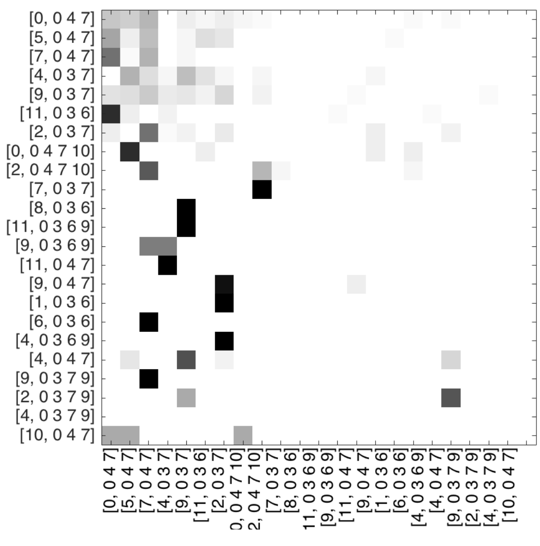

Figure 4: The first-order Markov transition matrix of chords (GCT groups) in a major Bach Chorales dataset.

3.1 A Markov transition matrix that accommodates two harmonic spaces

Aim of the proposed methodology is to construct a *musically meaningful* matrix of GCT chord transitions that includes and *extends* the respective transition matrices of two *initial* idioms. Figure 5 illustrates the parts of the *extended transition matrix* of the C major and F♯ major initial idioms referred to in the introduction as a ‘toy’ example. This matrix is built around the transition matrices of the initial idioms, $\mathtt{I_1}$ (C major) and $\mathtt{I_2}$ (F♯ major), with new transitions being inserted embodying the blends that are generated by combining pairs of transitions belonging to the two initial idioms. The matrix in Figure 1 is obtained by the same processes as the one in Figure 5, but includes fewer blends because of different blending setup, as described in Section 3.4. Each initial idiom is considered to incorporate a separate set of GCT chords, even if some chords might have common names in both idioms. For instance, even though the [0, 0 4 7] GCT chord might be found both in the Bach Chorales and Jazz datasets, it is treated within each idiom as a different chord, since it has a potentially different functional role in terms of the chords that come before or after it in each dataset, i.e. in terms of the transitions involving it. However, identical or similar chords in the two initial idioms are seen as "natural" harmonic connection points between these idioms; transitions between such chords are constructed in a pre-blending stage, described later.

The extended transition matrix is structured in blocks identified as follows: $\mathtt{I_1}$ and $\mathtt{I_2}$ are the transition matrices of the initial idioms, $\mathtt{A_{i-j}}$ are blended transitions that lead from idiom $\mathtt{I_i}$ to idiom $\mathtt{I_j}$, $\mathtt{B_{i-X}}$ (and $\mathtt{B_{X-i}}$) are blended transitions that lead from idiom $\mathtt{I_i}$ to a new chord generated with transition blending (respectively from a new chord to idiom $\mathtt{I_i}$), and C represent transitions between two new chords (not considered in the current implementation). Regarding the graph-

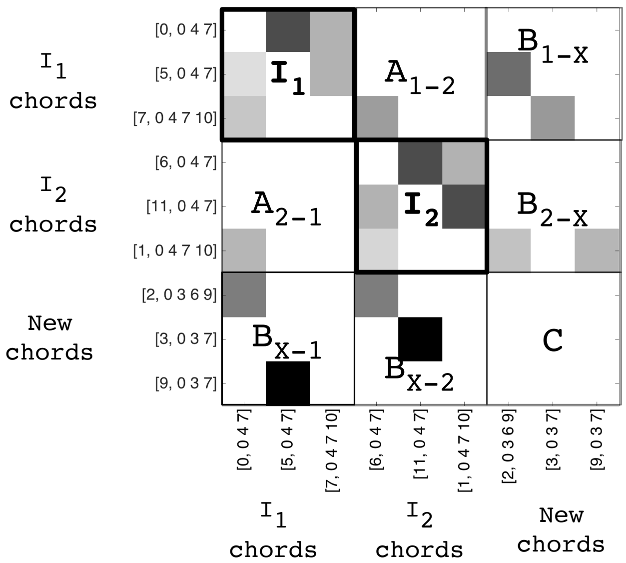

Figure 5: Graphical description of a *compound* matrix that includes transition probabilities of both initial idioms and of several new transitions generated through transition blending. These new transitions allow moving across the initial idioms, creating a new compound idiom.

ical representation of the extended matrix as depicted in Figure 5 the following facts should be highlighted:

1. By using transitions in $\mathtt{I_i}$, only chords of the i-th idiom are used. When using the transition probabilities in $\mathtt{I_i}$, the resulting harmonisations preserve the probabilistic character of idiom i.

2. Transitions in $\mathtt{A_{i-j}}$ enable direct jumps from chords of the i-th idiom to chords of the j-th idiom. If a blended transition happens to be in $\mathtt{A_{i-j}}$ there is no need for further considerations – such transitions can be directly included in the extended matrix.

3. Transitions in $\mathtt{B_{i-X}}$ go from a chord of idiom i to a new chord created with transition blending. Similarly, transitions in $\mathtt{B_{X-j}}$ 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 required: $c_i \rightarrow c_x$ followed by a transition $c_x \rightarrow c_j$, where c_i is in idiom i and c_j is in idiom j. A chain of two consecutive transitions with one intermediate external chord from chords of i to chords of j will be denoted as $\mathtt{B_{i-X-j}}$.

4. Transitions in C incorporate pairs of chords that are new to both the i-th and the j-th idioms. Having two external chords, transition blends in C are disregarded in the present work and, therefore, all probabilities in C are set to 0.

The proposed process for constructing the compound matrix intuitively generates new transitions by blending the most common transitions in the initial ($\mathtt{I_i}, i = 1, 2$) Markov matrices. It is not straightforward how a blended transition should be inserted in the extended matrix, since the extended matrix is a means to *interconnect and relate* chords between $\mathtt{I_i}$ and $\mathtt{I_j}$. The idea behind the proposed methodology is that blended transitions should allow moving from chords in $\mathtt{I_i}$ to ones in $\mathtt{I_j}$ and vice-versa. However, since transition blending potentially invents transitions between 'new' chords not belonging to any of the initial idioms, for the proposed method we ensure that every transition should have at least one chord that departs from or leads to $\mathtt{I_i}, i = 1, 2$ (and so should have at most one new chord created through blending), discarding transitions between two newly created chords. Therefore, in this study we postulate that blended transitions should include only one new 'intermediate' chord for moving from $\mathtt{I_i}$ to $\mathtt{I_j}$, and so we need to *ensure* that if a new chord is used, it will be preceded by a chord in $\mathtt{I_i}$ and followed by a chord in $\mathtt{I_j}$. If this requirement is not met, the new chord would be either a terminal or a beginning chord, constituting a 'dead-end' or 'unreachable' chord in the Markov chain.

3.2 Connecting transition tables via common chords

Two harmonic spaces may share common chords, or *similar* chords, which are hereby defined as chords that belong to the same GCT chord group. According to the methodology presented in Kaliakatsos-Papakostas et al. (2015), two chords belong to the same group (and are thus considered similar) if they (i) have the same root; (ii) have subset-related chord types; and (iii) both include pitch classes that are diatonic to the scale of the idiom. For instance, in a C major key, the GCT-represented chords [0, 0 4], [0, 0 4 7] and [0, 0 4 7 11] belong to the same major tonic group, while [0, 0 4 7 10] belongs to another since the pitch class value 10 is not diatonic to the major scale (this chord is a secondary dominant seventh to F major). For the remaining of this section, the term 'similar chords' will be used for describing chords that belong to the same GCT group.

The first step for generating the compound version of two transition matrices does not include blended transitions, but transitions that are composed of identical or similar chords between the two initial spaces – formulating an initial set of $\mathtt{A_{1-2}}$ and $\mathtt{A_{2-1}}$ transitions. These transitions allow moving between the two initial spaces by using common or similar chords as harmonic connection points. To this end, all possible transitions of such chords (i.e. all preceding and next chords) in one input idiom $\mathtt{I_i}$, are also considered as possible transitions of this chord in the other input idiom $\mathtt{I_j}$, "activating" the respective transitions in $\mathtt{A_{1-2}}$ and $\mathtt{A_{2-1}}$.

An example of this process is illustrated in Figure 6, where the similar chords of two simple 3-chord C major and D major spaces are combined. In this example, the chords G7 in C major and G in D major ([7, 0 4 7 10] and [7, 0 4 7] in GCT representation respectively) are similar in the context of C major. Since these chords are similar in at least one context (C major), the following transitions are activated:

1. All transitions leading to [7, 0 4 7 10] in C major (C2 column on top-left) should be also leading to [7, 0 4 7] (C2 column on top-right).

2. All transitions departing from [7, 0 4 7 10] in C major (R1 row on top-left) should be also departing from [7, 0 4 7] (R1 row on bottom-left).

3. All transitions leading to [7, 0 4 7] in D major (C1 column on bottom-right) should be also leading to [7, 0 4 7 10] (C1 column on bottom-left).

4. All transitions departing from [7, 0 4 7] in D major (R2 row on bottom-right) should be also departing from [7, 0 4 7 10] (R2 row on top-right).

The probability values assigned to the copied rows and columns that correspond to the similar chords are finally defined according to the *probability intensity multiplier* (PIM) described later in Section 3.4.

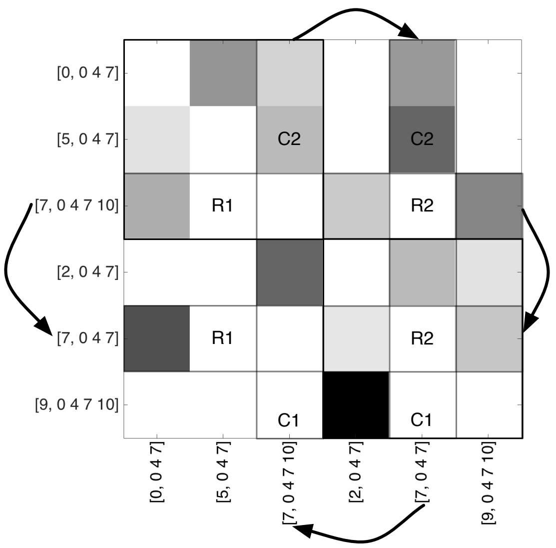

Figure 6: Graphical representation of shared transitions between similar chords between C major (top left) and D major (bottom right) spaces, where chords [7, 0 4 7 10] and [7, 0 4 7] are similar in the context of C major.

3.3 Identifying meaningful transition blending candidates

In order to reduce the number of applications of chord transition blending, the 10 most common transitions in $\mathtt{I}_1$ and $\mathtt{I}_2$ are gathered in two chord transition sets that represent the respective initial idioms. Every transition of idiom 1 is blended with the ones of idiom 2, producing 100 different potential applications of blending. Some applications of blending, however, may potentially subsume others, in a sense that some transition blending pairs may incorporate harmonic characteristics that have already been examined in other pairs. For the current study, meaningful transition blends are considered the ones that incorporate *maximal* subsets of features from the generic spaces in regards to the *subsumption* relation, as explained in the next paragraphs.

Each pair of input chord transitions (x_1, x_2) defines a generic space ($\mathcal{G}_{(x_1,x_2)}$) and a set of all possible generated blended transitions ($\mathcal{B}_{(x_1,x_2)}$); the generic space represents the common properties of the two input transitions, as described in Section 2.2, and by extension it defines the set of all possible blended transitions that fulfil its requirements, generated by Algorithm 1 and later ranked as per Section 2.3. It should be reminded that in the proposed transition blending methodology, only the properties related to pitch classes are considered in the generic space, namely the *fromPCs* and *toPCs* properties.

A generic space ψ_1 is said to *subsume* a generic space ψ_2 (or ψ_2 is subsumed by ψ_1), denoted as $\psi_1 \sqsubseteq \psi_2$, if ψ_1 is more general than or equal to ψ_2 (or equivalently ψ_2 is more specific than or equal to ψ_1) in the sense that ψ_1 defines a larger set of possible generated blended transitions. Using the above notation, $\mathcal{G}_{(x_1,x_2)} \sqsubseteq \mathcal{G}_{(y_1,y_2)}$ is equivalent to $\mathcal{B}_{(x_1,x_2)} \supseteq \mathcal{B}_{(y_1,y_2)}$, which means that the pair of transitions (y_1, y_2) gives rise to a smaller set of blended transitions $\mathcal{B}_{(y_1,y_2)}$ due to a more constrained generic space (i.e. the generic space $\mathcal{G}_{(y_1,y_2)}$ is more specific than $\mathcal{G}_{(x_1,x_2)}$).

The subsumption relation between generic spaces defines a *partial order relation*, that is, in the set of all possible generic spaces the subsumption relation satisfies, for all ψ_i, ψ_j, the following properties:

1. reflexivity: $\psi_i \sqsubseteq \psi_i$;
2. antisymmetry: if $\psi_i \sqsubseteq \psi_j$ and $\psi_i \sqsupseteq \psi_j$ then $\psi_i = \psi_j$; and
3. transitivity: if $\psi_i \sqsubseteq \psi_j$ and $\psi_j \sqsubseteq \psi_k$ then $\psi_i \sqsubseteq \psi_k$.

Therefore, in every finite subset of generic spaces there is at least one *maximal element*[7] ψ_M, for which no other generic space ψ_i is more specific than ψ_M. In other words, in any set of pair of transitions there are always (one or more) pairs that define maximal (i.e. most specific) generic spaces, while other non-maximal pairs are characterised by weaker conditions (i.e. with less strict generic spaces) that produce larger blending sets containing those produced by the most specific generic spaces.

At this point it is helpful to consider the notion of a 'blending quadruple', denoted by $\mathcal{Q}$, which is a set that comprises four elements: input transition 1 (I_1), input transition 2 (I_2), their generic space ($\mathcal{G}$) and a list of all the produced blended transitions ($\mathcal{B}$). Therefore, the blending quadruple that corresponds to the i-th pair of input transitions is a set described as: $\mathcal{Q}_i = \{I_1^i, I_2^i, \mathcal{G}_i, \mathcal{B}_i\}$. The subsumption relation is used on the generic spaces of quadruples to discard the ones that are 'overshadowed' by others, i.e. that incorporate a larger number of comparable common properties between the inputs. As an example, table 3 shows two quadruples that take part in the blending process with the initial spaces C major and F♯ major of the 'toy' example above. Thereby it can be observed that $\mathcal{G}_2 \sqsubseteq \mathcal{G}_1$ and therefore $\mathcal{G}_2$ (and the blends it includes) is not regarded in the subsequent processes. Formally, by considering the set $\mathcal{G} = \{\mathcal{G}_i\}$ of generic spaces from all blending quadruples that are generated for two initial idioms, the blends that are retained are the ones that correspond to maximal generic spaces in $\mathcal{G}$. The generic spaces that are completely empty are not considered in this process, since they subsume any other generic space and would be entirely disregarded. Since each blended transition in every $\mathcal{B}_i$ also has a rating value, this set can be further reduced by applying a threshold on rating value or on the number of desired outputs. For the next steps, a maximum of 100 blends with highest rating values will be retained for further processing in each $\mathcal{B}_i$.

For computational efficiency, in the implemented system the quadruple rejection/acceptance step precedes the blending step – therefore the blends of the discarded quadruples are actually never computed. In the case where two blending quadruples include identical generic spaces with different input transitions, the blends of both quadruples are retained. Even if the blends in quadruples with identical generic spaces are the same, they are evaluated differently, since the input transitions that produced them are different. Therefore, these quadruples include blends that are ranked differently, leading to different selections of topmost blends in the subsequent steps.

Table 3: Two example blending quadruples obtained by combining C major with F♯ major spaces.

Inputs		$\mathcal{G}$			$\mathcal{B}$
I_1^i	I_2^i		*fromPCs*	*toPCs*	
I_1^1: [7, 0 4 7 10] → [0, 0 4 7]	I_2^1: [1, 0 4 7 10] → [6, 0 4 7]	$\mathcal{G}_1$:	{{5, 11}	*ANY*}	$\mathcal{B}_1$
I_1^2: [7, 0 4 7 10] → [0, 0 4 7]	I_2^2: [11, 0 4 7] → [1, 0 4 7 10]	$\mathcal{G}_2$:	{{11}	*ANY*}	$\mathcal{B}_2$

3.4 Assigning probabilities and embedding transition blends in the extended idiom matrix

For each quadruple that passes the generic space subsumption filtering stage, the topmost 100 blends are kept while the rest are discarded. For each of these blends, a probability value is calculated, that will be used in subsequent selection steps described later. The proposed approach for assigning

[7] ψ_M is maximal with respect to $\{\psi_1, \psi_2, \ldots, \psi_n\}$ when there is no element ψ_i with $\psi_M \sqsubseteq \psi_i$, or equivalently, when for each i it either holds that $\psi_i \sqsubseteq \psi_M$ or ($\psi_M \not\sqsubseteq \psi_i$ and $\psi_i \not\sqsubseteq \psi_M$).

probability values to the blends of a quadruple is intended to reflect (a) the probability values of the input transitions that produced these blends and (b) the ranking placement of each blend in the blending quadruple. Specifically, if the probability value (in the initial transition matrix of the idiom) of the inputs that produced a blend is p_{I_1} and p_{I_2}, then the *potential* of a blend, p_b, is computed as:

$$p_b = \frac{p_{I_1} + p_{I_2}}{2} \, \frac{\text{rate}(b)}{\text{rate}_{\text{max}}},$$

where rate(b) is the rating value of the blend and rate_{max} is the maximum rating value in the examined blending quadruple. The final probability value of each blend is computed by normalising each row of the matrix to sum to one, according to the potential of each blend. In other words, the final probability assigned to a blend is the mean probability of the inputs that produced it, scaled by a factor that indicates the rating of this blend in comparison to the best-rated blend that these inputs have produced – the better the rate of the blend, the closer its probability value to the mean value of probabilities of the inputs.

Among the blending quadruples that are preserved, a number of their best blends is stored for further investigation, creating a pool of best blends. Based on trials, a large number of the best blends (i.e. 100) from each blending quadruple should be inserted in the pool of best blends, so that several scenarios for connecting the initial spaces can be created, since a greater number of blends in the pool of best blends introduces a larger number of possible commuting paths in $\mathtt{A_{i-j}}$ or in $\mathtt{B_{i-X-j}}$.

Blended transitions in the pool of best blends are, then, categorised according to whether they belong to category $\mathtt{A_{i-j}}$, $\mathtt{B_{i-X}}$ or $\mathtt{B_{X-i}}$. Blends that belong to category $\mathtt{A_{i-j}}$ can be directly embedded in the extended transition matrix. However, blends that belong to either $\mathtt{B_{i-X}}$ or $\mathtt{B_{X-i}}$ may potentially constitute terminal or beginning transitions respectively, as discussed in Section 3.1. Therefore, blends in $\mathtt{B_{i-X}}$ or $\mathtt{B_{X-i}}$ are matched in $\mathtt{B_{i-X-j}}$ chains/pairs and considered as integrated elements. The rating value assigned to every chain of blended transitions is the mean of ratings of each blend in the chain.

For allowing different intensities of blending in the harmonisations that the system produces, there are also two parameters, namely *rating-based selection* (RBS) and *probability intensity multiplier* (PIM), that define the number of blends to be embedded in the extended matrix and the relative value of probabilities of transitions outside the initial harmonic spaces ($\mathtt{I_1}$ and $\mathtt{I_2}$). RBS is in the range $[0, 1]$ and defines the percentage of top blends or transition chains that are imported in the extended matrix. For instance a RBS value of 0.5 imports 50% of the most highly rated blends, while a value of 0 generates an extended matrix that includes only the initial spaces and the pre-blending common/similar connections (see Section 3.2). For instance, the RBS value in the transition matrix of Figure 1 is at 0.03, while a greater RBS value (0.06) is used for constructing the matrix in Figure 5, which includes more blended transitions. PIM is in the $[0, 1]$ range and is used as a multiplier of all probabilities outside the $\mathtt{I_1}$ and $\mathtt{I_2}$ according to the following formula:

$$p_{\text{new}} = (0.1 + 20PIM)\ (p_{\text{old}}^{(1-\frac{PIM}{2})}),$$

where p_{new} is a new value (potentially greater than 1) assigned in the transitions matrix in the place of p_{old}, which is its probability value assigned by either the pre-blending or the blending stage. After all new values have been assigned, the transition matrix is normalised so that every row adds to 1. A PIM value of 0 reduces the probabilities of transitions in $\mathtt{A_{i-j}}$, $\mathtt{B_{i-X}}$ and $\mathtt{B_{X-i}}$, resulting in harmonisations that most probably use transitions (and, subsequently, chords) exclusively from one initial harmonic space, making connections between initial spaces less likely. Conversely, a PIM value of 1 increases the probabilities in $\mathtt{A_{i-j}}$, $\mathtt{B_{i-X}}$ and $\mathtt{B_{X-i}}$, encouraging harmonisations that transit from $\mathtt{I_i}$ to $\mathtt{I_j}$ and vice-versa by using existing chords (transitions in $\mathtt{A_{i-j}}$) or even new ones (transitions in $\mathtt{B_{i-X}}$ and $\mathtt{B_{X-i}}$), producing results that incorporate mixed harmonies of the initial spaces, as well as new chords.

4 Examples of blended harmonisations and methods for evaluation

Evaluating creative systems is a difficult task since there is no well-established and commonly accepted definition of creativity (e.g. Boden (2004); Wiggins (2006); for a comprehensive discussion see Jordanous (2013), chapter 3). The methods proposed so far either focus on the creative processes *per se* (Colton et al., 2011), or on the products of creative processes (Ritchie, 2007; Jordanous, 2013). Additionally, concept invention through blending harmonies is a new, currently vague and unexplored field. When blending harmonic styles there is no concrete expectation about the result. There are, however, specific problems in harmony that need to be resolved creatively (i.e. in novel and consistent ways), in which tools like concept invention can be used to propose many alternative and diverse solutions within a unified framework. A brief overview of some empirical evaluation techniques that have been employed (Zacharakis et al., 2017) or are part of ongoing work are presented as well (incorporating the 'raw' output produced by the assistant).

It is reminded that the presented methodology focuses on conceptual blending at the level of chord transitions (chord types consisting of pitch classes), not on voice leading and chord layout rules. The proposed harmonisation assistant employs elementary voice-leading algorithms (primarily for the outer voices); if, however, the goal is to reach a 'correct' harmonisation according to the voice leading rules of a specific musical style, additional human intervention is required. In the examples presented below, we have manually altered the voice leading provided by the assistant where necessary to make the examples more legible, while the 'raw' output of the system is given in the Appendix.

4.1 Example harmonisations

In these examples, we investigate whether the products of the system[8] are within an acceptable range of musical solutions, taking into account their aesthetic contexts. To the best of our knowledge, there is no methodology for resolving creatively harmonic 'problems' similar to the ones examined in this paper. Seven different musical idioms and some of their blends were used for the harmonisation of the melodies used in the presented examples, appearing in the following list:

- BC major & BC minor: The homophonic tonal harmonic idiom of J. S. Bach Chorales.
- JA major & JA minor: Mainstream jazz harmony, as encountered in selected jazz standards from the *Real Book*.
- CN: Greek composer Yannis Constantinidis's 20th-century modal style, as encountered in his '44 Greek miniatures for piano' (Tsougras, 2010).
- HM: Paul Hindemith's 20th-century neo-tonal harmonic idiom, as expressed in his 'Six Chansons' for a capella choir.
- WT: Whole-Tone post-tonal harmony, as encountered in selected excerpts from early 20th-century works.

An emphasis was given on the use of tonal idioms (mainly BC major and minor) in the present paper, although the dataset[9] used includes several diverse harmonic idioms, because major-minor tonality is probably one of the most studied harmonic idioms, so it functions as a reference point for testing and demonstrating blending procedures. The system produced raw MIDI files that were further processed by humans using music notation software (Finale 2014). The process involved two stages: 1) correction of music notation issues and enharmonic spellings of pitches in the MIDI file, and 2) manual editing of the produced harmonisation regarding separation of the bass line in

[8] For all the examples that follow, if not otherwise explicitly stated, the PIM and RBS values in the presented harmonisations are 0.7 and 0.9 respectively for the 'high blending parameters' setup and 0.3 and 0.5 respectively for the 'low blending parameters' setup.

[9] `http://ccm.web.auth.gr/datasetdescription.html`

a different layer, preservation of a constant number of active voices in the musical texture through the use of octave doublings, and reworking of the inner voices for smoother voice-leading where needed (although a strict application of common-practice voice-leading rules was not pursued). Also, manual analysis of harmonic progressions through the use of Latin roman numeral notation of tonal harmonisation was made in certain cases. The pitch content of the chords was always kept intact, and the bass line was manually altered in very few cases (indicated by * followed by the parenthesized original bass note in the scores) in order to improve the readability of the examples and put the focus on the blending results rather than the voice leading; the Appendix includes the 'raw' output format as produced by the system of all the examples. In summary, the examples illustrate harmonisations in which the pitch-class content of the chords (chord types) and almost all of the pitches of the bass line are computed automatically by the system, whereas inner voice leading is partly determined by hand.

4.1.1 Blending major and minor tonalities

The first harmonic setting that can be addressed with the proposed methodology concerns the problem of harmonising a melody with a blended major and minor harmonic style. For this purpose, a short 8-bar melody was constructed that deliberately avoids the 3rd and 6th melodic degrees of the C scale, making it ambiguous regarding its classification as major or minor. It consists of two 4-bar phrases (half cadence - full cadence) that form an 8-bar period. The short melody is harmonised four times (see Figures 7 and 12). It should be noted that apparent parallelisms between the input melody (upper staff) and some voices of the produced harmonisation (lower staff) are not considered errors here, because they merely reflect the property that melody notes (input) are always included in the harmony (output), regardless of voice assignment. Further separation of these harmonic 'blocks' into coherent voices following specific voice-leading practices might be handled in a post-processing stage, but this was not considered here. The first harmonisation is based on the BC major dataset and uses mainly the I and V harmonic degrees, with sparse use of subdominant-function chords (IV and ii). The second harmonisation uses the BC minor dataset and produces similar results in C minor (use of i, V and iv) with the exception of the final tonic chord, which is major, as encountered in the Bach chorales. The first blended harmonisation (low blending parameters, Figure 7 (c)) includes a mix of the chords encountered in the two unblended versions (I, V, iv), avoiding the minor tonic (i) and major subdominant (IV) chords and introducing other chords not used previously (♭VI, v, ii^o). The second blended harmonisation (high blending parameters) uses a different mix of chords from the two parallel tonalities, that includes the minor tonic (i) and the submediant (vi).

4.1.2 Blending different major keys for modulation purposes

Compound harmonic spaces created through transition blending can also be used for key transpositions, by combining the spaces of the incorporated keys. An original purpose-made 10-bar melody, which includes tritonal shifts from C major to F♯ major and a return to C major, is harmonised three times (see Figures 8 and 13), all with the use of the BC major (Bach Chorales major) harmonic idiom. The first harmonisation does not incorporate any blending of keys, so the harmoniser attempts to assign chords without modulating away from C major. The result reveals that the melodic shift towards F♯ major has been ignored, however the system has managed to assign chords to the melody's chromatic pitches, albeit with functionally awkward or ambiguous results (in b. 6-8 G♯ has been harmonised with E major chord, F♯ with F♯ diminished chord, G♯ with G♯ diminished chord and A♯ enharmonically with B♭ major chord, see Figure 8 (a) for an analytical attempt with Roman numerals). The second harmonisation uses the blended C-F♯ major space with low PIM and RBS values, so the system is now able to identify the modulating segment of the melody, and manages to suggest functionally correct chords for both the shift towards F♯ and the return to C major, as the harmonic analysis reveals (see Figure 8 (b)). All the chords are triads, except for the dominant, which appears with its 7th in two cases. The third harmonisation uses the blended space again, but with high PIN and RBS values. The result is quite original, as apart from the modulation, which has been identified and harmonised correctly, the system introduces chromaticism within each tonal region, with unexpected assigned chords in several cases (e.g. the B minor 7th in b. 1, the D

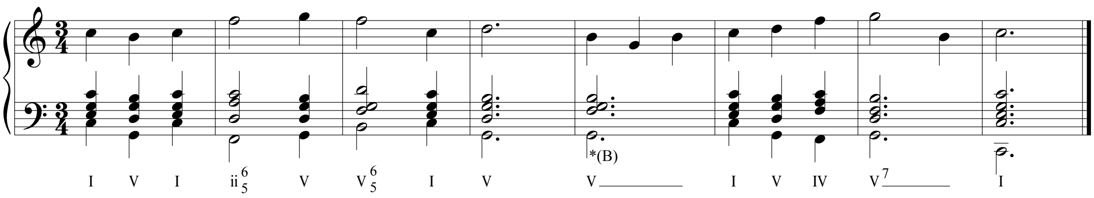

(a) Melody harmonised with the BC major idiom.

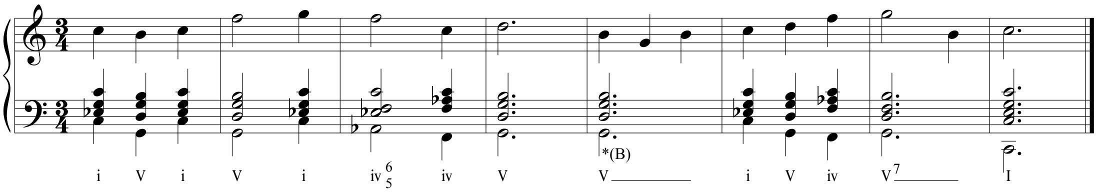

(b) Melody harmonised with the BC minor idiom.

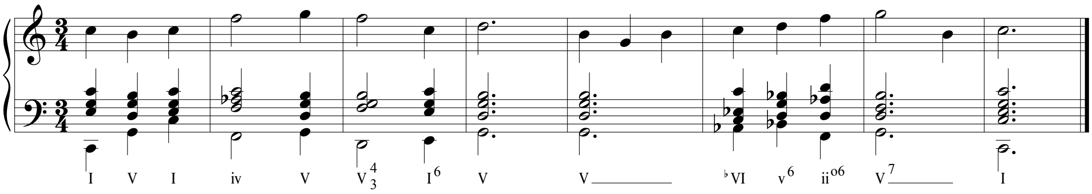

(c) Melody harmonised with the blended major-minor BC space, using low blending parameters.

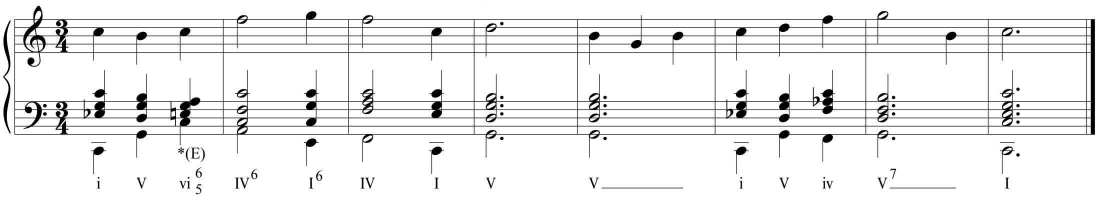

(d) Melody harmonised with the blended major-minor BC space, using high blending parameters.

Figure 7: A melody without the 3rd and 6th diatonic degrees harmonised with idioms learned from a dataset of (a) major (BC major), (b) minor (BC minor) Bach chorales and their harmonic blend with (c) low PIM and RBS values (0.05 and 0.1 respectively) and (d) high PIM and RBS values (0.7 and 0.9 respectively).

half-diminished in b. 6, the G minor 7th and C major 7th in b. 9). This harmonisation displays elements of unexpected originality to an extent that an explicit functional harmonic analysis would be unsuitable, and so has been avoided.

4.1.3 Creative blending of different major keys

Combining harmonic spaces in different keys can also be used creatively, e.g. for introducing chromatic harmonic characteristics when harmonising even simple tonal melodies, without necessarily focusing on the problem of key transposition. The Scottish folk song *'Ye Banks and Braes'* in the G major pentatonic mode can be harmonised in many different ways, with blended variations of the BC major harmonic idiom in various shifted tonalities (i.e. blended G major and other major tonalities). As an example, in Figures 9 (and 14) we present a harmonisation produced using the G major-B major blend, which features forced chromaticism applied on the pentatonic/diatonic space implied by the melody; the harmonisation with the plain BC major idiom (in G major) is a typical diatonic tonal harmonisation. Harmonic analysis has not been included in this example, because of instances where the harmony deviated from functional progressions towards free chromaticism,

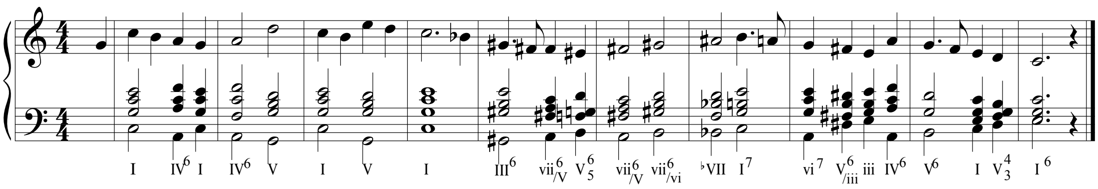

(a) Melody harmonised with the C major BC space.

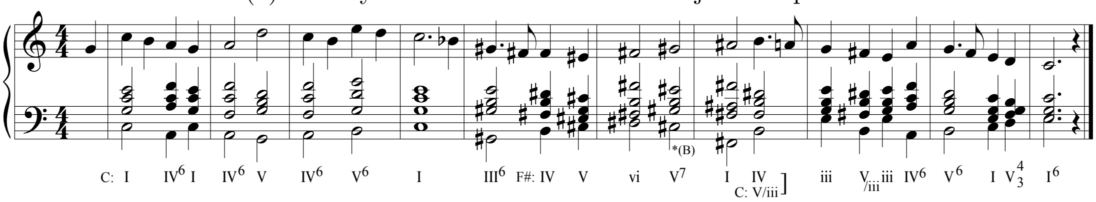

(b) Melody harmonised with the blended C and F♯ BC major spaces, using low blending parameters.

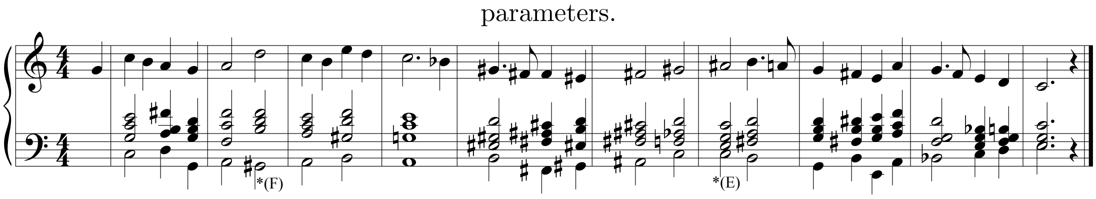

(c) Melody harmonised with the blended C and F♯ BC major spaces, using high blending parameters.

Figure 8: A melody harmonised with the learned BC major idiom with (a) C major tonality and blended versions of BC major in the tonality of C and its transposition in F♯ major with (b) low PIM and RBS values (0.05 and 0.1 respectively) and (c) high PIM and RBS values (0.7 and 0.9 respectively).

although such an analysis could be obtained. Many interesting harmonic phenomena potentially result out of such harmonisations, as, for instance, the augmented 6th chord in b. 5 and the chromatic linear progressions in b. 6-7 and b. 11 in the presented example. One interesting fact is that the produced chords cannot always be explicitly identified as belonging to one of the blended spaces. What is equally interesting is the fact that the blended tonal spaces can produce such a diverse range of forced harmonic chromaticism, with elements of tonal mixture, chords of ambiguous functionality, and chromatic contrapuntal chords even though the melody is purely pentatonic and does not even hint at chromaticism.

Figure 9: A traditional Scottish melody harmonised with the learned BC major space in G major and B major.

4.1.4 Creative blending of diverse harmonic spaces

Blending can be used creatively for combining two idioms from different eras. As an example, the '*Ode to joy*' melodic theme from L. v. Beethoven's 9th Symphony (transposed into C major) is harmonised three times in the sequel (see Figures 10 and 15). The first harmonisation uses the BC major idiom (no blending) and consists of the alternation of only two triads: the tonic and the dominant in root position and without 7th extensions (Figure 10 (a)). The second uses the JA idiom (no blending also) and conforms to the mainstream jazz harmony rules: all chords include major or minor 7ths, the main harmonic pattern is the ii-V-I turnaround and there is a tonicisation of the IV at the beginning of the second phrase instead of a half cadence (Figure 10 (b)). The third harmonisation is based on the harmonic idiom produced by blending the BC and JA idioms. As shown in Figure 10 (c), there is a mix of simple and extended triads and a combination of tonicisations and progressions in the circle of 5ths. Interestingly, now the tonicisation of IV occurs in b. 2-4 through a turnaround, a tonicisation of ii is prepared but avoided in b. 5, and a chromatic tonicisation of vi is observed in b. 7. These elements were not part of the unblended versions and seem implicitly only related to either idiom, although there is a sense that the jazz idiom dominates the system's choices.

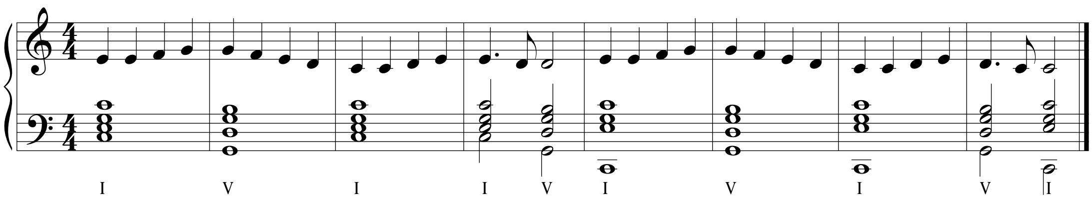

(a) The melody of *Ode to joy* harmonised with the BC major space.

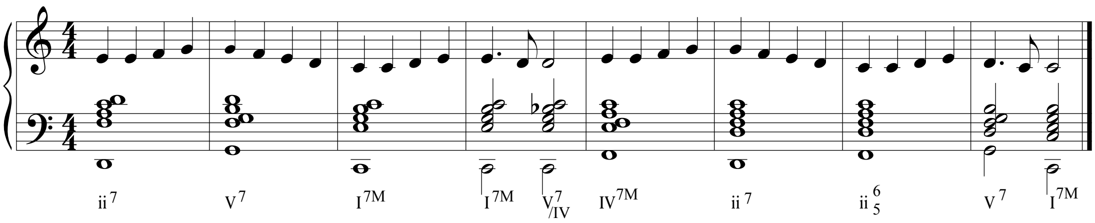

(b) The melody of *Ode to joy* harmonised with the JA major space.

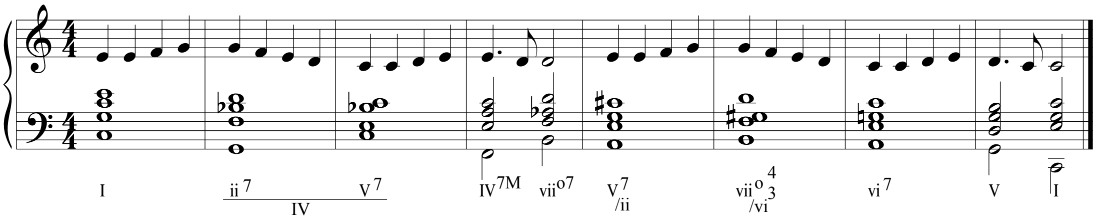

(c) The melody of *Ode to joy* harmonised with the the blended Bach chorale and jazz standards styles.

Figure 10: Beethoven's *Ode to Joy* theme harmonised in the style of (a) Bach chorales, (b) jazz standards and (c) their blended harmonic space.

Interesting harmonisations can also be produced with this methodology when harmonies of diverse and potentially idiosyncratic idioms are blended. Two example harmonisations are presented, involving the Greek folk melody *'Apopse ta mesanychta'* (Tonight at midnight) in D Aeolian mode. These examples are illustrated in Figures 11 and 16, using blends of diverse, mainly post-tonal, harmonic idioms. The first harmonisation uses a blend of Yannis Constantinidis's harmonic style (20th-century chromatic modal harmony, see Tsougras (2010)) and Whole-Tone harmony. It seems that Constantinidis's harmony dominates, with its parallel voice-leading (b. 1, 3), free use of minor or major 7th chords, and conclusion on an open-5th sonority (b. 8), however a characteristic influence of the WT space can be observed in b. 5 (WT chord C-D-E-G♯). The second harmonisation is based on a blending between minor jazz harmony (extended tonal idiom) and the neo-tonal harmonic idiom of Hindemith (free chromatic harmony with pitch centres). The chords suggested by the system are mainly extended triads with loose harmonic functions, but there are some notable exceptions, either as free mildly dissonant chords (mostly free verticalisations of diatonic sets), e.g. the A-D-G-B sonority in b. 2 and 4, and the quartal chords D-E-A-B (b. 5) and C-F-G-B♭ (b. 6) or as the highly dissonant sonority B-E♭-G-B♭ in b. 8. However, although certain elements of the harmonisations may be classified as stemming from one of the blended idioms, the overall produced idioms feel original and coherent.

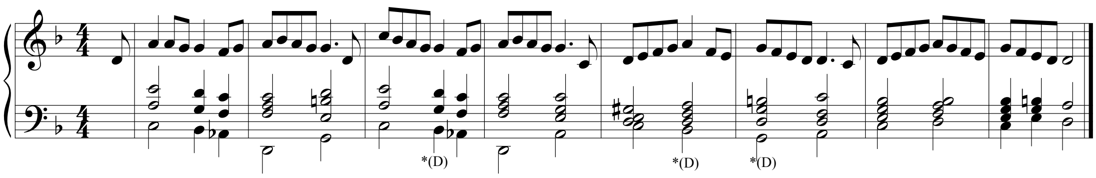

(a) A traditional Greek melody harmonised in the harmonic style of CN blended with the WT harmony.

(b) A traditional Greek melody harmonised in the harmonic style of JA blended with the HM harmonic idiom.

Figure 11: A traditional Greek melody harmonised with the blended harmonic styles of (a) Constandinidis with whole-tone and (b) jazz standards with Hindemith-style harmony.

4.2 Empirical evaluation methods

This section gives a brief overview of the so far employed and ongoing empirical evaluation methods for the produced outcome of the CHAMELEON melodic harmonisation assistant, which utilises the proposed methodology. First, the results obtained and presented in (Zacharakis et al., 2017) are shortly described, and then, some pointers to ongoing research are given. A detailed description of methods for the empirical experiments and the analysis of their results are beyond the scope of this paper. The empirical evaluation methods address two main questions:

1. How are harmonisations from blended idioms perceived by musically trained listeners?

2. Can the results produced by the system enhance human creativity or assist composers in exploring new harmonic areas?

Results obtained from empirical experiments on the raw output of the assistant indicate that the proposed blending methodology produces harmonies that are perceived as belonging to hybrid or new styles that potentially convey characteristics from both blended styles, while at the same time are useful as a basis of inspiration for composers.

The first question above regarding perception of blended idioms involves two experiments in a perceptual study presented in Zacharakis et al. (2017): a) mode and idiom and b) type of chromaticism classification of harmonisations. For mode classification, five harmonisations in major, minor and blended styles (BC major, BC minor and three blends with different blending parameter values) of a melody similar to the one presented in Figure 12 were presented simultaneously to 40 students from the Department of Music Studies at the Aristotle University of Thessaloniki. During the listening process, the participants were asked to rate the harmonisations on a Likert scale with 5 points from minor to major and a sixth option indicated as 'Other'. The results indicated that the participants clearly classified the pure major and minor harmonisations to the correct mode, while the blended harmonisations were mainly placed in the middle, with a small percentage categorising them as 'Other'. In a similar setting harmonisations of five melodies (two tonal, two jazz and a modal Greek traditional melody) were provided, in the harmonic styles learned from sets of Bach chorales (tonal) and jazz standards, along with some rather different harmonisations using models learned from the Beatles and Hindemith idioms in order to introduce diversity in the experiments. The five melodies were harmonised with the non-blended versions of the input idioms as well as with blended versions. As in the case of the major-minor mode classification experiment, participants were asked to categorise each harmonisation to the 'Tonal' or 'Jazz' categories, using a Likert scale from 1 to 5, or to categorise them as 'Other'. The results show that the participants were able to identify non-blended tonal from jazz harmonisations, while the blends where mostly categorised as in-between tonal and jazz or as belonging to the category 'Other'. Preference ratings, that were also requested in the study, indicated that the blended harmonisations of the system were equally preferred along with the non-blended ones; i.e. listeners think that the aesthetic value of the blended harmonisation is at least as good as the original harmonisation.

The type of chromaticism classification was examined in a study with 30 participants, all students from the Department of Music Studies at the Aristotle University of Thessaloniki. Participants were asked to identify whether harmonisations of the 'Ye Banks and Braes' melody (presented in Section 4.1.3) in G major alone and with different blended tonality combinations of the BC major idiom belong to the 'Diatonic', 'Chromatic', 'Atonal' or 'Other' categories. Additionally, participants were asked to rate each stimulus according to preference and expectancy (as opposed to novelty) within a range from 1 to 5. While the harmonisation obtained by the non-blended BC major idiom was unanimously categorised as 'Diatonic', the blended variations were mainly placed in the 'Chromatic' category. The expectancy of the blended harmonisations was significantly decreased in comparison to the non-blended one, while at the same time the preference ratings were alike for non-blended and blended harmonisations. Therefore, using the system in this way allows the generation of harmonisations that are perceived as unexpected chromatic and at the same time preferred, by blending a diatonic idiom (that is otherwise only able to produce expected diatonic harmonisations) with a transposed version of itself.

Regarding the question of whether the results produced by the proposed methodology could enhance human creativity or assist composers in exploring new harmonic ideas, ongoing research focuses on two empirical methods. In the first one students from the Department of Music Studies at the Aristotle University of Thessaloniki are asked to harmonise a melody before and after they listen to some harmonisations produced by the CHAMELEON system. More specifically, the melody of 'Ye Banks and Braes' is used (because of its clear tonal character) as the melody to be harmonised. The students are first asked to freely harmonise this melody according to their personal taste – not necessarily complying with any tonal rules of harmony. After a period of two weeks, they hand their first attempt on harmonising this melody. Then they are provided with some harmonisations of the same melody that were produced by the CHAMELEON melodic harmonisation assistant using blended idioms between transposed versions of the BC major style (a list of 6 harmonisations including the blended harmonisation presented in Section 4.1.3 as well as harmonisations in other modulation combinations). At this point, students are instructed to listen to those harmonisations and decide if they want to change something in their initial harmonisation and if so, make a new version. Preliminary results obtained so far among a small number of participants, indicate that almost everyone decided to change their initial harmonisation significantly, going from harmonisations that were closer to the tonal idiom to ones that incorporated more chromaticism. These results, however, are still at a preliminary stage and the collection of data from more participants is necessary to allow more conclusive results.

Our second approach to studying the enhancement of human musical creativity involves composers. The ongoing study examines how the output of CHAMELEON can be used creatively by composers as a source of harmonisation ideas that can provide novel harmonic frameworks/backgrounds for free development of musical textures. The task given to the composers involves the elaboration and variation of a melody of their choice for the composition of a piano miniature. In a first experiment four traditional melodies were given, along with a wide variety of diverse harmonisations produced by a system using the blending methodology. The composers selected the harmonisation/s of their melody that they considered most interesting/inspiring, and created their miniature compositions. Seven such miniatures for piano were performed by Fani Karagianni in a concert held at the Macedonian Museum of Contemporary Art in Thessaloniki on 19 October 2016 [10]. Even though the musical results of this attempt might arguably be considered as interesting and successful per se, discourse analysis (Wodak and Meyer, 2009) on interviews with the composers will be used in an effort to capture more detailed and robust indications about the actual influence that composers acknowledge regarding the output of the transition blending methodology.

5 Conclusions

Creative melodic harmonisation through the combination of harmonic spaces is examined in this paper; conceptual blending is integrated into a Markov chord transition methodology that is part of an idiom-independent melodic harmonisation assistant that learns from harmonic data. The algorithmic framework for conceptual blending developed in the context of the COINVENT (Concept Invention Theory) project is utilised to blend transitions of chord (pairs of successive chords), referred to as the 'input transitions', between different harmonic idioms, producing new blended transitions. Transition blending is based on combining a priori defined transition features, while the best blends are identified through a process that ranks all resulting blends according to the salience of their incorporated features; the salience of features is automatically induced by statistical assessment on the learned input idioms. The best blends of the most usual transitions in two initial harmonic idioms are then used to construct a new 'compound' harmonic space, that includes chords and observed transitions of the initial idioms along with new chords and new transitions created by transition blending.

The creative harmonic capabilities of the system have been examined under several melodic harmonisation blending settings, revealing a number of different previously unexplored cases (in the context of musical artificial intelligence) where this methodology is applicable – from robust problem solving to purely experimental harmonic exploration. The examples presented in this paper discuss some potentially interesting applications of the proposed methodology for blending: (i) major and minor modes; (ii) different keys for modulations; (iii) different keys for increasing chromaticism; and (iv) different, potentially diverse harmonic styles for exploring novel harmonic ideas. Results illustrate on the one hand that the presented system can be used for generating conventional harmonic solutions, constituting a useful tool e.g. for music education, and on the other hand, that it could be used for generating unconventional harmonic material, thus providing musicians with a tool that produces many creative alternatives in harmonising a given melody. This methodology is also within reach for potentially interested amateur musicians who might want to experiment with different idiom blending combinations, and it could be easily integrated in other software (e.g. games) for producing novel, unique and diverse musical backgrounds. The new possibilities that the proposed system offers, as algorithmically approached in the context of the COINVENT project, highlight the new capabilities that are introduced in computational creativity by conceptual blending. Those capabilities can be further enhanced in future work by developing more sophisticated probabilistic or rule-based voice leading methodologies that allow the generation of 'raw' output that adheres to the voicing layout forms of a given style more persuasively.

Evaluation of how the products of this system are perceived has been performed in (Zacharakis et al., 2017); results indicate that the intended purposes of blending are met, with the system creating compound idioms that are perceived either as in between the blended ones, or as something completely new. Additionally, pilot results indicate that the creativity of human composers is enhanced when they are assigned with harmonic composition tasks that involve the results of the proposed methodology; ongoing research in this direction will hopefully allow the extraction of more

[10] http://ccm.web.auth.gr/creativeusecomposers.html

conclusive results. An interesting future research direction is towards increasing the self-awareness of the system, by developing methods that automatically categorise the products of the system, either by performing style classification or qualitative evaluation. Additionally, the statistical modelling of learning different harmonic modules (i.e., chord types, transitions, cadences and bass voice leading) can be used per se for style identification or for harmonic similarity, either between blended and non-blended harmonisations or between harmonisations in corpora of different styles. It would be also interesting to examine how current state-of-the-art algorithms for style classification would classify blended harmonisations (as belonging to one of the input spaces, in-between them, or as belonging to a whole new style). Self-awareness on this level would allow the system to re-adjust the blending parameters, i.e. PIM and RBS, so that more meaningful blended harmonisations are produced without user intervention.

Acknowledgements

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

Appendix

This section includes the non-modified output of all the examples presented in Section 4 as produced by the system.

Figure 12: 'Raw' system output of all the examples in Figure 7.

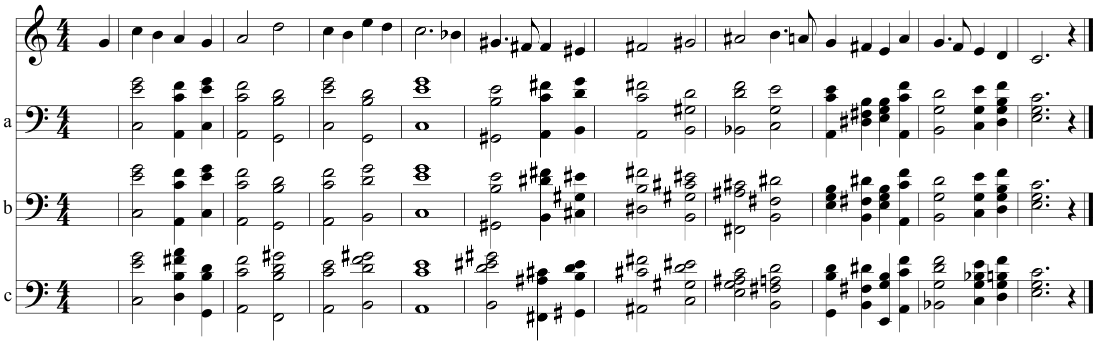

Figure 13: 'Raw' system output of all the examples in Figure 8.

Figure 14: 'Raw' system output of the example in Figure 9.

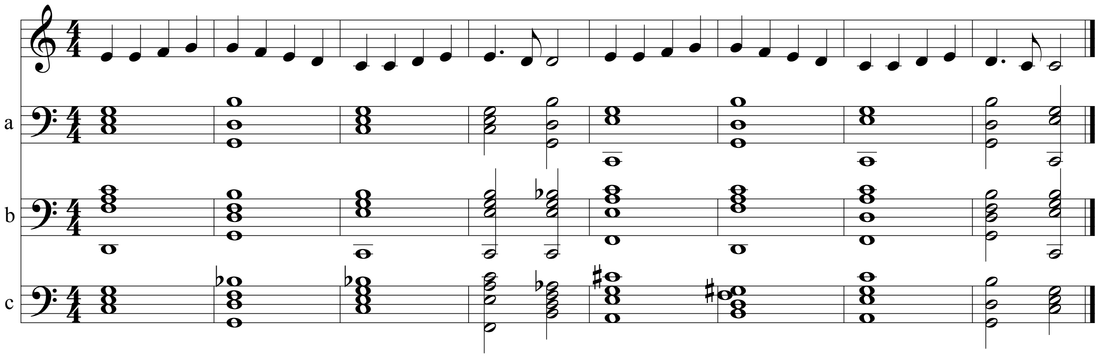

Figure 15: 'Raw' system output of all the examples in Figure 10.

Figure 16: 'Raw' system output of all the examples in Figure 11.

References

Antovic, M. (2011). Musical metaphor revisited: Primitives, universals and conceptual blending. *Universals and Conceptual Blending (February 17, 2011)*.

Boden, M. A. (2004). *The creative mind: Myths and mechanisms*. Psychology Press.

Boden, M. A. (2009). Computer Models of Creativity. *AI Magazine*, 30(3):23.

Callender, C., Quinn, I., and Tymoczko, D. (2008). Generalized voice-leading spaces. *Science*, 320(5874):346–348.

Cambouropoulos, E. (2012). A Directional Interval Class Representation of Chord Transitions. In *In Proceedings of the Joint Conference ICMPC-ESCOM 2012 (12th International Conference for Music Perception and Cognition, & 8th Conference of the European Society for the Cognitive Sciences of Music)*, ICMPC-ESCOM 2012.

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., Kaliakatsos-Papakostas, M., and Tsougras, C. (2015). Structural blending of harmonic spaces: a computational approach. In *Proceedings of the 9th Triennial Conference of the European Society for the Cognitive Science of Music (ESCOM)*.

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.

Chordia, P., Sastry, A., Mallikarjuna, T., and Albin, A. (2010). Multiple viewpoints modeling of tabla sequences. In *ISMIR*, volume 2010.

Cleary, J. G. and Teahan, W. (1995). Experiments on the zero frequency problem. In *Proc. Data Compression Conference*, volume 480.

Colton, S., Pease, A., and Charnley, J. (2011). Computational Creativity Theory: The FACE and IDEA Descriptive Models. In *2nd International Conference on Computational Creativity*.

Confalonier, R., Schorlemmer, M., Plaza, E., Eppe, M., Kutz, O., and Peñaloza, R. (2015). 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*.

Cook, N. (2001). Theorizing musical meaning. *Music Theory Spectrum*, 23(2):170–195.

Dixon, S., Mauch, M., and Anglade, A. (2010). Probabilistic and logic-based modelling of harmony. In *Exploring Music Contents*, pages 1–19. Springer.

Ebcioglu, K. (1988). An expert system for harmonizing four-part chorales. *Computer Music Journal*, 12(3):43–51.

Eppe, M., Confalonier, 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., editors, *Logic Programming and Nonmonotonic Reasoning*, volume 9345 of *Lecture Notes in Computer Science*, pages 309–316. Springer International Publishing.

Fauconnier, G. and Turner, M. (2003). *The Way We Think: Conceptual Blending And The Mind's Hidden Complexities*. Basic Books, New York, reprint edition edition.

Goguen, J. (2006). Mathematical Models of Cognitive Space and Time. In Andler, D., Ogawa, Y., Okada, M., and Watanabe, S., editors, *Reasoning and Cognition*, volume 2 of *Interdisciplinary Conference Series on Reasoning Studies*. Keio University Press.

Goguen, J. and Harrell, D. F. (2010). Style: A Computational and Conceptual Blending-Based Approach. In Argamon, S. and Dubnov, S., editors, *The Structure of Style: Algorithmic Approaches to Understanding Manner and Meaning*, pages 147–170. Springer, Berlin.

Granroth-Wilding, M. and Steedman, M. (2014). A robust parser-interpreter for jazz chord sequences. *Journal of New Music Research*, 0(0):1–20.

Hirata, K., Tojo, S., and Hamanaka, M. (2014). Algebraic mozart by tree synthesis. In *Proceeding of the joint 11th Sound and Music Computing Conference (SMC) and 40th International Computer Music Conference (ICMC)*, ICMC–SMC 2014.

Jordanous, A. K. (2013). *Evaluating computational creativity: a standardised procedure for evaluating creative systems and its application.* PhD thesis, University of Sussex.

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., Cambouropoulos, E., Kühnberger, K.-U., Kutz, O., and Smaill, A. (2014). Concept Invention and Music: Creating Novel Harmonies via Conceptual Blending. In *In Proceedings of the 9th Conference on Interdisciplinary Musicology (CIM2014)*, CIM2014.

Kaliakatsos-Papakostas, M., Confalonier, R., Corneli, J., Zacharakis, A., and Cambouropoulos, E. (2016a). An argument-based creative assistant for harmonic blending. In *Proceedings of the 7th International Conference on Computational Creativity (ICCC).*

Kaliakatsos-Papakostas, M., Makris, D., Tsougras, C., and Cambouropoulos, E. (2016b). Learning and creating novel harmonies in diverse musical idioms: An adaptive modular melodic harmonisation system. *Journal of Creative Music Systems*, 1(1).

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)*, Malaga, Spain.

Koops, H. V., Magalhães, J. P., and de Haas, W. B. (2013). A functional approach to automatic melody harmonisation. In *Proceedings of the First ACM SIGPLAN Workshop on Functional Art, Music, Modeling & Design*, FARM '13, pages 47–58, New York, NY, USA. ACM.

Makris, D., Kaliakatsos-Papakostas, M., and Cambouropoulos, E. (2015). A probabilistic approach to determining bass voice leading in melodic harmonisation. In *Mathematics and Computation in Music: Proceedings of the 5th International Conference, MCM 2015*, London, UK. Springer, Berlin.

Moore, A. F. (2013). *Song means: Analysing and interpreting recorded popular song.* Ashgate Publishing, Ltd.

Ontañón, S. and Plaza, E. (2012). On knowledge transfer in case-based inference. In Agudo, Belén, D. and Watson, I., editors, *Case-Based Reasoning Research and Development*, volume 7466 of *Lecture Notes in Computer Science*, pages 312–326. Springer Berlin Heidelberg.

Ox, J. (2014). Analogy and conceptual blending are part of a visualization toolkit for artists and scientists: Introducing the cognitive space transfer. *Methods*, 1:6.

Pachet, F., Roy, P., and Barbieri, G. (2011). Finite-length markov processes with constraints. *transition*, 6(1/3).

Phon-Amnuaisuk, S., Smaill, A., and Wiggins, G. (2006). Chorale harmonization: A view from a search control perspective. *Journal of New Music Research*, 35(4):279–305.

Raphael, C. and Stoddard, J. (2004). Functional harmonic analysis using probabilistic models. *Computer Music Journal*, 28(3):45–52.

Ritchie, G. (2007). Some empirical criteria for attributing creativity to a computer program. *Minds and Machines*, 17(1):67–99.

Rohrmeier, M. (2011). Towards a generative syntax of tonal harmony. *Journal of Mathematics and Music*, 5(1):35–53.

Schorlemmer, M., Smaill, A., Kühnberger, K.-U., Kutz, O., Colton, S., Cambouropoulos, E., and Pease, A. (2014). Coinvent: Towards a computational concept invention theory. In *5th International Conference on Computational Creativity (ICCC) 2014*.

Spitzer, M. (2004). *Metaphor and musical thought*. University of Chicago Press.

Stefanou, D. and Cambouropoulos, E. (2015). Enriching the blend: Creative extensions to conceptual blending in music. In *Proceedings of the 9th Triennial Conference of the European Society for the Cognitive Science of Music (ESCOM)*.

Tsougras, C. (2010). The Application of GTTM on 20th Century Modal Music: Research Based on the Analysis of Yannis Constantinidis's "44 Greek Miniatures for Piano". *Musicae Scientiae*, 14(1 suppl):157–194.

Tsougras, C. and Stefanou, D. (2015). Conceptual Blending and Meaning Construction: A Structural/hermeneutical Analysis of the 'Old Castle' from Musorgsky's 'Pictures at an Exhibition'. In *Proceedings of the 9th Triennial Conference of the European Society for the Cognitive Science of Music (ESCOM)*.

Tymoczko, D. (2006). The geometry of musical chords. *Science*, 313(5783):72–74.

Whorley, R. P., Wiggins, G. A., Rhodes, C., and Pearce, M. T. (2013). Multiple viewpoint systems: Time complexity and the construction of domains for complex musical viewpoints in the harmonization problem. *Journal of New Music Research*, 42(3):237–266.

Wiggins, G. A. (2006). A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems*, 19(7):449–458.

Wodak, R. and Meyer, M. (2009). *Methods for critical discourse analysis*. Sage.

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)*, Malaga, Spain.

Zacharakis, A., Kaliakatsos-Papakostas, M., Tsougras, C., and Campouropoulos, E. (2017). Empirical methods for evaluating music structural blends: a case study on melodic harmonisation. *Musicae Scientiae*, page (forthcoming).

Zbikowski, L. M. (2002). *Conceptualizing Music: Cognitive Structure, Theory, and Analysis*. Oxford University Press.

Zbikowski, L. M. (2008). Metaphor and music. *The Cambridge handbook of metaphor and thought*, pages 502–524.