

CONCEPT INVENTION AND MUSIC: CREATING NOVEL HARMONIES VIA CONCEPTUAL BLENDING

Maximos Kaliakatsos-Papakostas¹, Emiliios Cambouropoulos¹,
Kai-Uwe Kühnberger², Oliver Kutz³ and Alan Smaill⁴

¹ School of Music Studies, Aristotle University of Thessaloniki, Greece

² Institute of Cognitive Science, University of Osnabrück, Germany

³ Institute of Knowledge and Language Engineering, Otto-von-Guericke University Magdeburg, Germany

⁴ School of Informatics, The University of Edinburgh, UK

Correspondence should be addressed to: maxk@mus.auth.gr

Abstract: Conceptual blending is a cognitive theory whereby elements from diverse, but structurally-related, mental spaces are “blended” giving rise to new conceptual spaces that often possess new powerful interpretative properties, allowing better understanding of known concepts or the emergence of novel concepts altogether. This paper provides an overview of the wide computational methodological spectrum that is being developed towards building an automatic melodic harmonization system that employs conceptual blending, yielding harmonizations that inherit characteristics from multiple idioms. Examples of conceptual blending in harmony are presented that exhibit the effectiveness of the developed model in inventing novel harmonic concepts. These examples discuss the invention of well-known jazz cadences through blending the underlying concepts of classical music cadences, as well as the construction of larger chord sequences. Furthermore, examples of a conceptual blending interpretation in human compositions that motivated the goals of the system’s design are given. Finally, conceptual blending between harmonic and non-harmonic domains is discussed, offering tools that allow for intuitive human intervention in the harmonization process.

1. INTRODUCTION

New concepts may be invented by “exploring” previously unexplored regions of a given conceptual space (exploratory creativity) or transforming in novel ways established concepts (transformational creativity) or by making associations between conceptual spaces that were previously not directly linked (combinational creativity); Boden maintains that the latter, i.e., combinational creativity, has proved to be the hardest to describe formally [1].

Conceptual blending is a cognitive theory developed by Fauconier and Turner [2, 3] whereby elements from diverse, but structurally-related, mental spaces are “blended” giving rise to new conceptual spaces that often possess new powerful interpretative properties allowing better understanding of known concepts or the emergence of novel concepts altogether. Conceptual blending is a process that allows the construction of meaning by correlating elements and structures of diverse conceptual spaces. It relates directly to Boden’s notion of combinational creativity. In the context of the European FP7 project COINVENT¹, outlined in [4], a formal model for conceptual blending has been detailed based on Goguen’s initial ideas of a Unified Concept Theory [5]. It incorporates important interdisciplinary research advances from cognitive science, artificial intelligence, formal methods and computational creativity. To substantiate the model’s potential, a proof-of-concept autonomous computational creative system that performs melodic harmonization is being developed. Furthermore, some non-musical examples of concept invention through the COINVENT methodology have been recently demonstrated, towards establishing the model’s functional characteristics. These examples discuss the creation of novel monster-like entities through blending ontologies of animals [6], as well as the invention of the complex numbers by utilizing simpler mathematical concepts [7]. The general computational workflow has been described in some detail in [8].

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 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, which fosters the selection and combination of concepts that inject novelty and creativity to the melodic harmonization process. A rich idiom-independent representation of harmonic concepts is proposed: from the “primitive” chord events, see General Chord Type (GCT) representation [9], to a hierarchical multiple-viewpoint representation of harmonic structure that allows “meaningful” blends at various hierarchic levels of harmony. Alongside the development of an algorithmic part that facilitates conceptual blending, a large dataset of more than 400 harmonically annotated pieces from various diverse musical idioms is being constructed; the knowledge extracted by these annotated pieces (mainly in a statistical formalization) will comprise the rich background required for interesting and creative blends.

In the context of the COINVENT project, a core-model for conceptual blending has been developed, shown in Figure 1. According to this model, conceptual blending is employed between two input spaces, I_1 and I_2 , yielding a “blendoid”², i.e. a conceptual space that includes a blend of concepts from both input spaces. Often, the blendoid will not be created directly from I_1 and I_2 , but rather is based on generalizations of some concepts that pertain either to I_1 or I_2 , leading to the respective “weakened” input spaces, I_1^* and I_2^* . ‘Weakening’ can be based on a variety of techniques, and its purpose is to remove irrelevant information, and to avoid obvious and ‘uninteresting’ clashes resp. formal inconsistencies. According to [11], for instance, a term A may be considered to be more general than a term B if everything that is described by B is also described by A . For example, the term “a red European minivan” is more general than the term “a red German minivan”. As it is intuitively expected, two conceptual spaces, I_1 and I_2 , will most likely include concepts/terms that are incompatible and/or contradicting. To this end, a generalization operation employed for the construction of the weakened input spaces, aims to resolve these contradictions, allowing the compatible and non-contradicting parts of the input spaces to be combined/blended. One particular method that is explored within COINVENT is to compute such weakenings based on the idea of “amalgams” [12].

However, when considering conceptual blending for input spaces based on rich knowledge repositories, there are numerous possibilities for generalization, leading to questions about which generalization scenario would yield a “useful” blend. Thereby, the utilization of a rich background is required for setting some rational, field-specific limitations to the generalization possibilities that would potentially make sense, as well as a mechanism for

¹www.coinvent-project.eu

²The notion of blendoid was introduced by Goguen, see e.g. [10], and is further discussed technically in [4, 8].

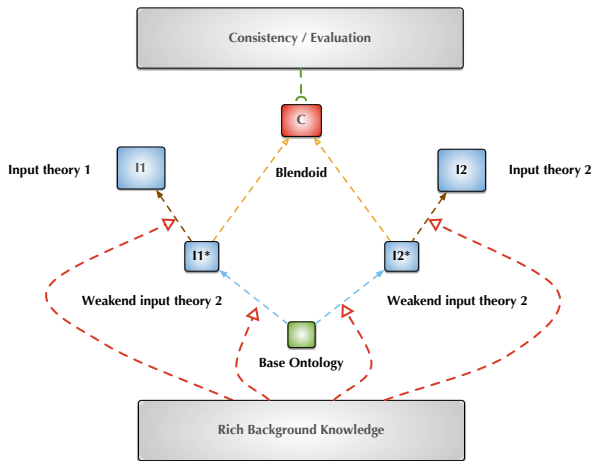


Figure 1: The COINVENT core model for conceptual blending between two input spaces (I_1 and I_2).

consistency/evaluation check of the resulting-blended conceptual space. Especially for music generation, the evaluation process regarding conceptual blending should be tackled with “objectivity”; although some may argue that evaluation should apply to art in a more “relaxed” manner, the proposed evaluation mainly concern the demand to have blended spaces that encompass characteristic attributes of the input spaces. A more extensive and thorough description of the subjects discussed in the last two paragraphs is beyond the scope of this paper, which emphasizes applications of conceptual blending that exemplify the hitherto discussed framework. For further details regarding evaluation in artificially intelligent systems, the interested reader is referred to [13], and for a more specific discussion regarding blending to [8].

A central aim of this paper is to provide an overview of the wide methodological spectrum that the development of the COINVENT melodic harmonizer introduces, as well as to present examples of conceptual blending in musical harmony. The following harmonic domains will be of particular interest:

1. *Chord-level:* Individual chords that share common properties are blended, giving rise to novel instances of chord functions (e.g. cadence blending).
2. *Chord-sequence level:* sequences of chords from different idioms are blended, creating sequences of different degrees of originality.
3. *Harmonic structure level:* different levels of harmonic structure from different harmonic spaces (i.e., different levels of induced probabilistic grammars) are combined, generating “coherent” new blends (e.g. phrase-level endings from one style combined with low-level chord transitions from another).
4. *Meta-level description:* Features characterizing a given harmony or melody are blended with features of a different harmonic space, allowing the generation of novel melodic harmonizations (such high-level features embody various structural properties and, additionally, extra-musical features such as mood, motion, tension, shape etc).

The presented examples, on the one hand, aim to demonstrate the effectiveness of the computational model w.r.t. “inventing” novel concepts (like the invention of the backdoor and the tritone substitution cadences in Section 2). On the other hand, some examples show how human compositions can be interpreted as conceptual blends, given in Section 4. These examples informed in particular the development of the methodological framework and general design of the COINVENT harmonizer.

Many aspects of the introduced examples, e.g. meta-level harmonic information discussed in Section 5, are part of ongoing research and entail challenging methodological questions that cannot be fully resolved in this paper and are left for future work.

2. CHORD-LEVEL BLENDING

Let us assume an early harmonic language in which mostly diatonic notes are allowed and dissonances in chords are essentially forbidden (except possibly using minor 7th intervals as in the dominant seventh chord). We assume that, in this early harmonic space, some basic cadences have been established as salient harmonic functions around which the harmonic language of the idiom(s) has been developed – for instance, the authentic/perfect cadence, half cadence, interrupted cadence and even older cadences such as the Phrygian cadence. The main question to be addressed in this section is the following: Is it possible to invent new cadences based on these basic cadence types? Is it possible for a computational system to create novel cadential schemata that function musically and cognitively as harmonic endings of musical phrases?

Such cadential harmonic transitions may be described formally as a set of chord primitives and transitions between them (see Table 1). These definitions incorporate only the very essential features of the cadences, i.e. chord type, root, bass (in a given context of key). If a cadence is represented in such a plain and relatively rigid fashion in a computational harmonic system, then only limited variation is enabled. Can such definitions be used for new cadence invention?

If exploration of the given space of, say, the authentic cadence description, is allowed, then modifications and alterations of the initial concept are possible; but are they meaningful? Various options of exploring the given definition may give many new chord progressions almost all of which will not be cadential and will not embody the dominant quality; for instance, if the root of the semi-final chord is replaced (say, by $X + 4$ or $X + 5$ or $X + 11$ etc.) or the *ch_type* (say, by minor or diminished chord type) then all sorts of likely or unlikely chord progressions (many containing non-diatonic notes) are generated, hardly any of which will embody the cadential character of the initial concept.

However, such a simple, if not naive, definition of authentic cadence may produce surprisingly interesting new concepts if combined with other cadential concepts via mechanisms of conceptual blending. Take, for instance, aligning and blending aspects of the phrygian cadence with the authentic cadence (Figure 2). If root: $X + 10$ of phrygian replaces root = bass = $X + 7$ in the major seventh version of the authentic cadence, then the backdoor progression (\flat VII7-I of jazz harmony) emerges! If bass: $X + 1$ of phrygian replaces root = bass = $X + 7$ in the major seventh version of the authentic cadence, then the tritone substitution progression (\flat II7-I) of jazz harmony emerges! Roots and bass notes of chords are rather important, so, in this instance, importing such features from one cadence to another give rise to surprising new cadential concepts (that actually appeared in music history centuries after the two input cadences).

The exciting invention of new cadences in this case should not be overestimated. There are a number of drawbacks that are inherent in the above elementary representation of the input cadences (relative prominence of chord features such as leading note and tritone are not explicitly represented and this blending may not work with other cadences). More sophisticated representations are part of ongoing research in developing a reliable cadence blending formal model.

While a thorough and extensive description of the discussed cadence blending example is beyond the descriptive scope of this paper, an overview of the underlying COINVENT model-wise theoretic aspects is provided. According to the methodological framework that supports conceptual blending, as roughly described in Section 1, the input spaces of the cadence example are the Authentic cadence (e.g. I_1) and the Phrygian cadence (e.g. I_2). The construction of the weakened input spaces that allow for the discussed cadence blending results includes the generalization in some cadence attributes, in a way that the “union” of the weakened spaces do not include contradicting attributes.

For instance, the semi-final chord type in I_1 is a major or a dominant seventh (type $[0, 4, 7]$ or $[0, 4, 7, 10]$) while in I_2 it is a minor (type $[0, 3, 7]$); a straightforward union of I_1 and I_2 would yield a contradiction – among others – about the semi-final chord’s type - it cannot be major and minor at the same time. For both the backdoor and the tritone substitution cadence examples, the contradiction is resolved by generalizing the minor chord type to the more general

Authentic Cadence	Phrygian Cadence
<p>Tonal space: scale: 0,2,4,5,7,9,11 % major scale scale: 0,2,3,5,7,8,11 % harmonic minor tonic: X % $X \in \{0,1,2,\dots,11\}$</p>	<p>Modal space: scale: 0,1,3,5,7,8,10 % phrygian mode tonic: X % $X \in \{0,1,2,\dots,11\}$</p>
<p>semi-final-chord: (ch_type: [0,4,7] % major chord root = bass = X+7) % P5 above tonic OR (ch_type: [0,4,7,10] % major seventh root = bass = X+7) AND final chord: ch_type: [0,4,7] OR [0,3,7] root = bass = X</p>	<p>semi-final-chord: ch_type: [0,3,7] % minor chord root: X+10 % 2M below tonic bass: X+1 % 2m above tonic AND final chord: ch_type: [0,7] OR [0,3,7] root = bass = X</p>
<p>Tritone Substitution</p>	

Figure 2: A simple formal definition of the authentic and phrygian cadences (without rhythmic information); last row: the result of blending between properties of the semi-final chords (bass: $X + 1$ of phrygian replaces $\text{root} = \text{bass} = X + 7$ in the major seventh version of the authentic cadence).

concept of TYPE – which also includes major and dominant seventh chords. The fact that the I_2 input space was weakened in terms of its semi-final chord type, leads to obtaining blendoids that include a chord type that is obtained by I_1 . It should be noticed that the semi-final chord type could as well be generalized in I_2 , leading to blendoids that have minor semi-final chords.

Accordingly, there are numerous scenarios for which concept to generalize from which input space, a fact that fosters numerous blendoids to emerge from two input spaces like the ones of the cadence example. However, it is imminent that some blendoids may make more sense than others (e.g. be more functional or aesthetically pleasing). Defining which blendoids are better than others is a matter of evaluation, which is a part of the undergoing COINVENT research. Additionally, depending on the formalization of concepts in the input spaces, the blendoids may also incorporate inconsistencies that should be addressed by specific rules that are included to a set of background knowledge facts.

3. CHORD-SEQUENCE-LEVEL BLENDING

A classical theoretical basis for chord sequences in the tonal system are cadences. In classical music theory, cadences are often used to describe a balance between tensions in a harmonic progression and a resolution of these tensions. Examples of well-known cadences are perfect cadences, half cadences, plagal cadences etc. In music pieces, chord progressions are clearly not restricted to combinations of small sets of well-defined cadences, but exploit a richer repertoire of chords sequences. We consider in this section longer sequences of chord progressions like descending fifths, descending fourths (i.e. ascending fifths), descending thirds and other progressions. Here are some properties of such chord progressions:

- There are different ways for representing chord sequences. A simple formalism to encode crucial properties of chord sequences are feature structures, where features like *key*, *mode*, *function*, *inversion* etc. together with appropriate values can be used to describe chords. Sequences of chords can be modeled as sequences of feature structures. Feature structures are used as an important formalism in different disciplines, for example, very prominently in linguistics [14].
- Representing chords with features like *key*, *mode*, *function* etc. has its limitations. For example, certain chords occurring in

late romanticism and chords in certain types of jazz music or modern music are often rather hard, sometimes ambiguous, and sometimes even impossible to analyze with the described features.

- Basic cadences of classical music often do not involve features like sevenths, ninths, elevenths etc. Nevertheless, the representation of general chord progressions should allow the coding of such features.
- Chord progressions occurring in pieces of music are more general than basic cadences in the sense that various degrees of freedom allow in progressions to depart from the root key (e.g. in modulations), to ignore rules imposed by cadences (e.g. to make a piece more interesting) etc.

Blending chord sequences in this framework can be considered as blending the involved chords of the single sequences in order to achieve a new sequence. It turns out that there are many design choices in a blending process of general chord sequences. In order to keep the complexity of the representation as simple as possible, we exemplify the blending of general chord sequences in non-trivial examples, but we are using an extremely simplified notation of chord sequence, namely as a sequence of functions similar to the presentation in scores of popular music.

Example 1 Two input sequences together with the generalization and candidates for blended sequences is listed:

Input 1: $T \rightarrow T_p \rightarrow S_p \rightarrow D7 \rightarrow T$ (in C major: $C \rightarrow a \rightarrow d \rightarrow G7 \rightarrow C$)
 Input 2: $T \rightarrow DDD \rightarrow DD \rightarrow D \rightarrow T$ (in C major: $C \rightarrow A \rightarrow D \rightarrow G \rightarrow C$)
 Generalization: $T \rightarrow x \rightarrow y \rightarrow D \rightarrow T$
 Blend 1: $T \rightarrow T_p \rightarrow DD \rightarrow D \rightarrow T$ (in C major: $C \rightarrow a \rightarrow D \rightarrow G \rightarrow C$)
 Blend 2: $T \rightarrow DDD \rightarrow S_p \rightarrow D7 \rightarrow T$ (in C major: $C \rightarrow A \rightarrow d \rightarrow G7 \rightarrow C$)

Blending requires as a first step the computation of a generalisation. This can easily be done by first, aligning the given sequences according to heuristics and second, by introducing variables for all those values of features which do not coincide. The analogy engine Heuristic-Driven Theory Projection (HDTP) is used to compute such generalisations [15]. In our example, we collapsed the dominant and the dominant seventh chord (fourth chord in the sequences) in the generalisation to a dominant chord. The two blended candidates for possible chord sequences emerge naturally from the input, which are undoubtedly two interesting chord sequences.

Example 2 Again two input sequences are considered, but this time the sequences are rather complex and they differ in the length of the sequence. In the following, the first input sequence is the circle of fifths. This sequence is represented by iterated subdominant chords (double subdominant, triple subdominant etc.), respect. iterated dominant chords.

Input 1: $T \rightarrow S \rightarrow SS \rightarrow S3 \rightarrow S4 \rightarrow S5 \rightarrow S6 \rightarrow D5 \rightarrow D4 \rightarrow D3 \rightarrow DD \rightarrow D \rightarrow T$
 Input 1 (C maj): $C \rightarrow F \rightarrow B^b \rightarrow E^b \rightarrow A^b \rightarrow D^b \rightarrow G^b \rightarrow B \rightarrow E \rightarrow A \rightarrow D \rightarrow G \rightarrow C$
 Input 2: $T \rightarrow SS \rightarrow T \rightarrow D5 \rightarrow T_p \rightarrow SS \rightarrow D5 \rightarrow T$
 Input 2 (C maj): $C \rightarrow B^b \rightarrow C \rightarrow B \rightarrow a \rightarrow B^b \rightarrow B \rightarrow C$
 Generalization: $T \rightarrow x1 \rightarrow x2 \rightarrow x3 \rightarrow x4 \rightarrow x5 \rightarrow x6 \rightarrow T$
 Blend 1: $T \rightarrow S \rightarrow T \rightarrow S3 \rightarrow T_p \rightarrow SS \rightarrow D5 \rightarrow T$
 Blend 1 (C maj): $C \rightarrow F \rightarrow C \rightarrow E^b \rightarrow a \rightarrow B^b \rightarrow B \rightarrow C$
 Blend 2: $T \rightarrow SS \rightarrow SS \rightarrow S3 \rightarrow S4 \rightarrow SS \rightarrow S6 \rightarrow T$
 Blend 2 (C maj): $C \rightarrow B^b \rightarrow B^b \rightarrow E^b \rightarrow A^b \rightarrow B^b \rightarrow G^b \rightarrow C$

Blend 3: T → S → SS → D5 → S4 → S5 →
 D5 → T
 Blend 3 (C maj): C → F → B \flat → B → A \flat → D \flat →
 B → C

The generated candidates for chord sequences can be considered as rather interesting. In particular, there seems to be a trade-off between the degree of interestingness of sequences and more conservative sequences that stick tentatively closer to one of the original input sequences. This could be metaphorically described as a more progressive form of chord sequences versus a more conservative form of chord sequence.

4. HARMONIC-STRUCTURE-LEVEL BLENDING

Conceptual invention and blending is facilitated when a rich background of diverse concepts is available and when these concepts are structured in such ways that creative mappings are supported. Concepts should be rich and highly structured. Humans are very capable in “compressing” complex networks of atoms and relations in higher-level more compact concepts, and then using these new simpler concepts in other more complex tasks. When concept invention takes place, the rich structural networks that often lie dormant under “seemingly” simple concepts, get activated enabling meaningful mappings and productive blends. A computational system of (musical) creativity must have access to such rich underlying structural representations on various hierarchic levels.

It is maintained that a melodic harmonisation assistant that facilitates conceptual blending should allow a modular highly structured representation of harmonic concepts in an explicit manner at various hierarchic levels and parametric viewpoints. In this study, five constituent structural components of harmony are explicitly represented:

1. *Harmonic pitch space*: scales, pitch hierarchies in scales, consonance/dissonance, chord types (e.g. in GCT).
2. *Chord transitions*: Learning of chord transitions from corpus data in certain idiom/style (e.g. dominant is followed most commonly by tonic).
3. *Cadences*: Learning of chord transitions that end phrases at various hierarchic levels (e.g. for tonal music, perfect cadence for the highest level cadence, other types of cadences at various lower level structural boundaries, etc.).
4. *Modulations*: Changes of harmonic pitch spaces that characterize a certain style (e.g. neighboring/distant tonalities, density of modulations, etc.).
5. *Voice leading*: through layering of individual voices (e.g. parallel/similar/oblique/contrary motion, drone tones, repetition or “compulsory” motion of certain pitches, preparation/resolution of dissonance etc.).

In the current study, at the lowest level the GCT representation [9] has been utilized for automatically encoding chords in the context of a given pitch space and consonance/dissonance ordering of intervals. Then this extracted encoding is used for harmonic learning at various levels. In [16] a constrained HMM (cHMM) was introduced that combines a well-studied probabilistic methodology, namely, the hidden Markov model (HMM), with constraints that incorporate fixed beginning and/or ending chords and/or intermediate anchor chords. The HMM is a very thoroughly studied technique that allows melodic harmonization that reflects the characteristics of the idiom they have been trained on, while several HMM modifications for melodic harmonization have been proposed in the literature, like for instance [17, 18, 19] and [20]. However, among the disadvantages of the HMMs – which the cHMM in combination with higher harmonic level techniques tackle – is that they are not able to capture dependencies on a large time scale, depending on their order.³ The usual HMM order utilized in the literature is 1, or sometimes 2, thus allowing dependencies of chord pairs or triplets to be sufficiently encoded.

Larger-scale chord dependencies can be captured by techniques

that consider the hierarchic architecture of harmony. Regarding the COINVENT melodic harmonization framework, beginning or ending chords (cadences) and intermediate chords (e.g. relating to tonicisations/modulations/phrase endings) are represented using simple probabilistic grammars that capture harmonic dependences among distant events. There is a considerably extensive literature for techniques that utilize grammars for harmonization, among which [21, 22, 23, 24], however the grammatical rules employed are based on theoretic considerations of specific idioms, while the approach followed in COINVENT applies a probabilistic context. Additionally, the application of efficient voice leading is also tackled through a statistical learning technique, which encapsulates statistical information about pitch height contour relations between the constituent pitches of chords. The phrase structure grammars and the voice leading statistical learning methodology are parts of ongoing research; further details about the aforementioned ongoing research are not provided considering space limitations, as well as the overall descriptive scope of the paper.

Once structural characteristics of diverse harmonic idioms are induced in an explicit modular fashion, then blends can be created that combine different harmonic aspects from different harmonic spaces. For instance, modal chord transitions may be combined with tonal cadences (or the reverse), or, more “adventurous” arbitrary blends may be generated that combine, say, atonal chord transitions with tonal voice leading and jazz cadences (represented as atonal GCTs). Such harmonic blending can be found in music by composers of different periods (most notably in the 20th century). Such compositions constitute motivating examples that forged the COINVENT methodological approach to melodic harmonization; some exemplar compositions are shortly presented in the next paragraph, exhibiting application paradigms of conceptual blending in different hierarchic harmonic levels, namely the level of key modulation and the level of chord transitions. The harmony in these examples incorporate modulations between tonalities that divide the octave into equal parts – thus called symmetric (e.g. major thirds, minor thirds or major seconds), while their inner-key transitions and cadences are tonal.

The harmonic plan of the theme of Giant Steps by John Coltrane can be seen as a blend between the tonal concept of descending circle of fifths (tonal cadential pattern: V7-I or ii7-V7-I) with a non-diatonic concept emerging from the symmetrical division of the octave into three equal parts (major thirds). The tonalities of the theme are structured as: $B - G - E\flat, G - E\flat - B, E\flat - G - B - E\flat - [B]$. Each tonality is established via a tonal cadential pattern: $(II7-V7)B-(V7)G-(V7)E\flat, (II7-V7)G-(V7)E\flat-(V7)B, (II7-V7)E\flat-(II7-V7)G-(II7-V7)B-(II7-V7)E\flat-[(II7-V7)B]$. Schubert begins his Piano Sonata D. 850, III with a sequence of ascending minor thirds through the keys of C, E \flat , and G \flat (may be interpreted as a blend between standard tonal cadences and symmetrical subdivision of octave into four equal parts for the modulation keys). Beethoven’s Waldstein sonata, begins with modulations/tonicisations that are separated by tones (IV-V7-I cadential patterns on G, F, G, A, G of the marked C major key), which can be seen (even though this is far fetched) as a blending between tonal cadences and parts of the whole-tone scale. A simplified and intuitive version of the COINVENT framework that enables such blends is illustrated in Figure 3. Various other such blends can be invented by employing the methodological context described in Section 1 on a systematic representation of different pitch/harmonic elements of diverse harmonic spaces in an explicit manner, allowing creative combinations of compatible constituent harmonic components.

The type of structural harmonic blending suggested in this section applies also to blending between melody and harmony in the sense that any melody implies a set of harmonic features that could be utilized to enforce blending between different harmonic idioms, as a means of “selecting” the proper harmonic elements from the respective trained harmonic spaces that are able to provide a satisfactory harmonization for the melody at hand. Furthermore, a melody can be harmonized using the prevailing implied harmonic system or other harmonic spaces that are relatively close to the prevailing harmonic system or even more remote harmonic spaces or blends between different harmonic systems that are consistent with the given melody.

³The order of a HMM or a Markov model in general, is the number of consecutive past states that define each current state.

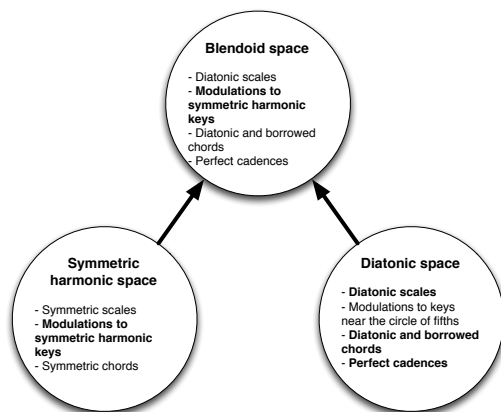


Figure 3: Simplified version of the blending diagram that produces the harmonic progressions in the Giant Steps by John Coltrane, Schubert’s Piano Sonata D. 850 and Beethoven’s Waldstein sonata.

5. META-LEVEL HARMONIC BLENDING

Humans commonly understand and describe aspects of music using conceptual models borrowed from other domains of human knowledge and experience. Analogy, metaphor, metonymy are mechanisms that enable such cross-domain mappings; the emerging conceptual models of music allow a richer understanding/experience of music. Concepts “borrowed” from other conceptual domains, such as concepts of shape (ascending-descending-concave curves), motion (parallel, oblique, contrary, acceleration, deceleration, static, dynamic), density (dense, sparse, thickening, thinning out), muscle-tone (tension-relaxation), balance (stability – instability), emotion/mood (happy, sad, nostalgic, melancholic, tranquil, sober, etc) can be “aligned” to music structural concepts, and be used to describe and understand better musical structures and processes. Metaphorical connections of concepts between different domains may lead to conceptual blending (see, for instance [25]).

A formal study regarding music and conceptual blending has been presented in [26], where a formal approach is presented that “blends” color spaces with musical chords. This study yielded some relation correspondences between colors and chords, however, as the authors suggest, the “fragility” of these blends allows only for domain-specific evaluation that incorporates direct human knowledge and intuition. This paper also mentions Mussorgsky’s “Pictures of an exhibition” as an example of artistic (not algorithmic) cross-domain blending in music; in this work, Mussorgsky creates “musical projections” of paintings, blending musical concepts with the ones embodied in the paintings. Additional examples of blending between musical and non-musical concepts have been presented in [27, 28], however, no formal mechanism that produces these blends is introduced.

Let us examine briefly the relation between affect and music structure. Emotions are often thought as having at least two qualities/dimensions, namely, valence that relates to pleasantness or hedonic value, and arousal that relates to bodily activity [29]. These two affective qualities are considered to be orthogonal to each other creating a two-dimensional space in which emotions can be placed (e.g. happiness has high arousal and positive valence whereas melancholy is low arousal and rather negative valence). The question that arises is: how is it possible for music to express feelings and emotions of listeners, or even to induce feelings/emotions in listeners? How can listeners perceive emotions/moods in abstract sound structures?

It is maintained in this paper that conceptual blending allows the alignment of two independent and unrelated input spaces, i.e. conceptual space of emotions and conceptual space of musical structure, in such a way a blended space between music and emotion arises (musical emotions have been studied extensively in recent years - see overview in [30]). The music emotion

blended space is to a large extent subjective, however, there are underlying principles that allow a certain amount of shared space in a certain cultural context. For instance, one possible such principle underlying valence (i.e. pleasantness) is the principle of simplicity: given certain sensory or other input, the simplest (most succinct) description is preferred and associated with a positive valence (cf. algorithmic theory of beauty by [31]). An important factor underlying arousal is bodily activity in terms of motion (from relaxation to high energy bodily motion) which can be described in more abstract terms as amount of energy/activity per time unit. Many music parameters/features can be rated in the continuum from simple to complex (e.g. diatonicism/chromaticism, consonance/dissonance, rhythmic/melodic pattern) or from low to high activity (e.g. tempo, harmonic rhythm, voice activity). Such music features can then be mapped onto the valence or arousal dimensions, and may give rise to particular music emotions (empirical studies have established relationships between structural properties of music with dimensions of emotions [32]).

Harmonic spaces of diverse harmonic systems can be described in terms of rich ontologies that contain not only music structural features (see previous section) but also non-structural features such as emotions/moods, types of motion, gestures, shapes, tension features and so on. All these will form ontologies consisting of complex networks of interrelated structural and non-structural features/properties that characterise harmonic spaces. Such harmonic spaces may be presented as input spaces to the COINVENT core model of conceptual blending, resulting in high-level blending of music harmonies. The description and construction of such ontologies is part of ongoing research.

6. CONCLUSIONS

We presented an overview of our ongoing research to tackle the problem of melody harmonization—the generation of harmonic accompaniments for a given melody—through a conceptual blending approach that can utilize harmonic concepts from multiple and diverse idioms. The algorithmic means towards this direction are provided by a computational model of creativity that is being developed in the context of the COINVENT project. This project aims to develop a computationally feasible, cognitively-inspired, formal model of concept invention, based on Fauconnier and Turner’s theory of conceptual blending. It is also based on a sound mathematical theory of concepts, following Goguen’s proposal of a Unified Concept Theory. Melody harmonization via conceptual blending is a challenging application area and allows for an in-depth evaluation of the discussed formal model. Moreover, the modeling challenges that arise lead to novel considerations concerning the entire automated harmonization edifice.

A first aim of our ongoing research is to approach conceptual blending in harmony through a systematic perspective, focusing on formalizing harmonic concepts, as well as observing the utilization of musical concepts by human composers. To this end, we presented and briefly discussed some initial examples concerning concept invention for cadences: the invention of the well-known tritone substitution and backdoor jazz cadences through blending concepts of the authentic and the phrygian cadences. Prototype implementations of the cadence blending are currently examined on the “*Ontohub*” (<https://ontohub.org>) web platform, where the COINVENT blending experiments will be hosted. Further, a conceptual blending interpretation of human compositions is given in specific examples, reflecting the rationale behind the selected methodological context, which incorporates capturing harmonic concepts through statistical learning on multiple hierarchic levels of harmony. Specifically, although an in-depth presentation of the algorithmic specifics is omitted, a combined scheme of probabilistic grammars and a modified version of the hidden Markov model (the constrained hidden Markov model) is developed. The goal of this methodological combination is to allow different learned concepts to be employed on different levels of the harmonisation process (e.g. cadences from one idiom combined with short-term chord progressions from another). Moreover, examples discuss the generation of entire chord progressions that incorporate blended characteristics from two input chord sequences. Finally, a more

intuitively-driven approach is explored by employing blending on a higher, meta-level description of harmony. Thereby, blending of more abstract harmonic concepts is allowed, leading to harmonisations that incorporate blending on a higher descriptive level.

Acknowledgements: The project COINVENT acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open grant number: 611553. The authors would like to thank Costas Tsougras and Nikos Diminakis for their valuable comments, especially regarding musical blending examples.

REFERENCES

- [1] M. A. Boden: *Computer Models of Creativity*. In *AI Magazine*, volume 30(3):23, 2009.
- [2] G. Fauconnier and M. Turner: *The Way We Think: Conceptual Blending And The Mind's Hidden Complexities*. Basic Books, New York, reprint edition, edition, 2003.
- [3] M. Turner: *The Origin of Ideas: Blending, Creativity, and the Human Spark*. Oxford University Press, 2014.
- [4] M. Schorlemmer, A. Smaill, K.-U. Kühnberger, O. Kutz, S. Colton, E. Cambouropoulos, and A. Pease: *COINVENT: Towards a Computational Concept Invention Theory*. In *5th International Conference on Computational Creativity (ICCC) 2014*. 2014.
- [5] J. Goguen: *Mathematical Models of Cognitive Space and Time*. In D. Andler, Y. Ogawa, M. Okada, and S. Watanabe (eds.), *Reasoning and Cognition, Interdisciplinary Conference Series on Reasoning Studies*, volume 2. Keio University Press, 2006.
- [6] F. Neuhaus, O. Kutz, M. Codescu, and T. Mossakowski: *Fabricating Monsters is Hard. Towards the Automation of Conceptual Blending*. In *Proceedings of the Workshop Computational Creativity, Concept Invention, and General Intelligence (C3GI 2014)*. 2014.
- [7] J. Fleuriot, E. Maclean, A. Smaill, and D. Winterstein: *Reinventing the Complex Numbers*. In *Proceedings of the Workshop Computational Creativity, Concept Invention, and General Intelligence (C3GI 2014)*. 2014.
- [8] O. Kutz, J. Bateman, F. Neuhaus, T. Mossakowski, and M. Bhatt: *E pluribus unum: Formalisation, Use-Cases, and Computational Support for Conceptual Blending*. In T. R. Besold, M. Schorlemmer, and A. Smaill (eds.), *Computational Creativity Research: Towards Creative Machines, Thinking Machines*. Atlantis/Springer, 2014.
- [9] E. Cambouropoulos, M. Kaliakatsos-Papakostas, and C. Tsougras: *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. 2014.
- [10] J. Goguen and D. F. Harrell: *Style: A Computational and Conceptual Blending-Based Approach*. In S. Argamon and S. Dubnov (eds.), *The Structure of Style: Algorithmic Approaches to Understanding Manner and Meaning*, pages 147–170. Springer, Berlin, 2010.
- [11] S. Ontañón and E. Plaza: *On Knowledge Transfer in Case-Based Inference*. In D. Agudo, Belén and I. Watson (eds.), *Case-Based Reasoning Research and Development, Lecture Notes in Computer Science*, volume 7466, pages 312–326. Springer Berlin Heidelberg, 2012.
- [12] S. Ontañón and E. Plaza: *Amalgams: A Formal Approach for Combining Multiple Case Solutions*. In *Proceedings of the 18th International Conference on Case-Based Reasoning Research and Development, ICCBR'10*, pages 257–271. Springer-Verlag, Berlin, Heidelberg, 2010.
- [13] S. Colton, A. Pease, and J. Charnley: *Computational Creativity Theory: The FACE and IDEA Descriptive Models*. In *2nd International Conference on Computational Creativity*. 2011.
- [14] G. Smolka: *Feature-constraint Logic for Unification Grammars*. In *Journal of Logic Programming*, volume 12(1-2):51–87, 1992.
- [15] M. Schmidt, U. Krumnack, H. Gust, and K.-U. Kühnberger: *Heuristic-Driven Theory Projection: An Overview*. In *Computational Approaches to Analogical Reasoning: Current Trends, Studies in Computational Intelligence*, volume 548, pages 163–194. Springer, 2014.
- [16] M. Kaliakatsos-Papakostas and E. Cambouropoulos: *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. 2014.
- [17] S. A. Raczyński, S. Fukayama, and E. Vincent: *Melody Harmonization With Interpolated Probabilistic Models*. In *Journal of New Music Research*, volume 42(3):223–235, 2013.
- [18] N. Borrel-Jensen and A. Hjortgaard Danielsen: *Computer-assisted music composition – A database-backed algorithmic composition system*. B.S. Thesis, Department of Computer Science, University of Copenhagen, Copenhagen, Denmark, 2010.
- [19] N. Yogev and A. Lerch: *A System for Automatic Audio Harmonization*, 2008.
- [20] I. Simon, D. Morris, and S. Basu: *MySong: Automatic Accompaniment Generation for Vocal Melodies*. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08*, pages 725–734. ACM, New York, NY, USA, 2008.
- [21] M. T. Granroth-Wilding: *Harmonic analysis of music using combinatory categorial grammar*. Ph.D. thesis, Institute for Language, Cognition and Computation School of Informatics University of Edinburgh, Edinburgh, Scotland, 2013.
- [22] M. Granroth-Wilding and M. Steedman: *A Robust Parser-Interpreter for Jazz Chord Sequences*. In *Journal of New Music Research*, volume 0(0):1–20, 2014.
- [23] M. Rohrmeier: *Towards a generative syntax of tonal harmony*. In *Journal of Mathematics and Music*, volume 5(1):35–53, 2011.
- [24] H. V. Koops, J. P. Magalhães, and W. B. de Haas: *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. ACM, New York, NY, USA, 2013.
- [25] M. T. Descamp: *Conceptual Blending And Metaphor Theory*. In *Metaphor and Ideology*, chapter 2, pages 19–62. Brill, 2007.
- [26] F. Camara Pereira and F. Amílcar Cardoso: *Knowledge Integration with Conceptual Blending*. In *Proceedings of the 12th Irish Conference on Artificial Intelligence & Cognitive Science*. Maynooth, Ireland, 2007.
- [27] L. M. Zbikowski: *Conceptualizing Music: Cognitive Structure, Theory, and Analysis*. Oxford University Press, 2002.
- [28] L. M. Zbikowski: *Metaphor and music theory: Reflections from cognitive science*. In *Music theory online*, volume 4(1):1–8, 1998.
- [29] L. F. Barrett: *Discrete emotions or dimensions? The role of valence focus and arousal focus*. In *Cognition & Emotion*, volume 12(4):579–599, 1998.
- [30] P. N. Juslin and J. Sloboda: *Music and Emotion*. In D. Deutsch (ed.), *The psychology of music*. Academic Press, 2012.
- [31] J. Schmidhuber: *Low-complexity art*. In *Leonardo*, pages 97–103, 1997.
- [32] P. Gomez and B. Danuser: *Relationships between musical structure and psychophysiological measures of emotion*. In *Emotion (Washington, D.C.)*, volume 7(2):377–387, 2007.