

A Transformational Creativity Tool to Support Musical Composition

Cristóbal RAYA ^{a,1}, Francisco J. RUIZ ^a, Cecilio ANGULO ^a, Albert SAMÀ ^a and Núria AGELL ^b

^a*BarcelonaTech, Vilanova i la Geltrú, Spain.*

^b*ESADE-URL, Sant Cugat, Spain*

Abstract. In this paper we use the idea of *conceptual space* introduced by Boden and redefine some properties such *appropriateness* and *relevance* that facilitate the computational implementation of the transformational creativity mechanism. While appropriateness can only be evaluated by an expert, relevance can be objectively measured for any spectator. Computational creativity is based on the relationship between appropriateness and relevance of a concept, and therefore a computational system can be used to support this task. The paper analyses this relationship in the field of music in order to obtain a computer tool to support the musical composition task.

Keywords. Computational creativity, Transformational creativity, Music composition

1. Introduction

For most of the history of Artificial Intelligence (AI), creativity has probably been the most difficult human faculty to replicate. This is mainly because creativity, like intelligence, is an elusive phenomenon to define. While AI studies the performance of tasks by means of computers or robots which would be deemed to require intelligence if performed by a human, Computational Creativity (CC) studies performances which would be deemed creative if performed by a human. Despite this parallelism, while AI has experienced remarkable advances in the last decades, CC is in an earlier phase of its development and it has only recently established as a research field with its own identity and agenda. Since 2010, the *Association for Computational Creativity* has organized an annual international conference in this field and numerous other computer science and AI conferences have also included sessions devoted to this area.

It is much easier to answer the question 'Where is creativity?' than to answer the question 'What is creativity?'. Creativity can be found in paintings, sculpture, literature, music, architecture, as well as in business, engineering, software development, scientific research, and almost all human activities. It is a phenomenon whereby something new and valuable is produced: such as an idea, a solution, a marketing strategy, a literary work, a painting, a cookery recipe or a musical composition. The difficulty of defining

¹Corresponding Author: BarcelonaTech, Av. Víctor Balaguer, s/n, 08800 Vilanova i la Geltrú, Spain; E-mail: cristobal.raya@upc.edu

creativity is clear from the number of definitions that can be found in the literature. Taylor [1] gives about 50 definitions of creativity. Authors have diverged in the precise definition of creativity beyond these two features: *originality* (new) and *appropriateness* (valuable).

As a research field, the goal of CC is threefold: to model, to simulate and to serve as a support tool for creative tasks. Formalizing models of creative behaviour to implement them on a computer enables us to gain a better understanding of the mechanism of thought and, therefore, it helps us to better understand ourselves. In addition, it also enables us to understand social and cultural interaction as creativity is an individual phenomenon, as well as a collective process.

The simulation of creative behaviour enables us to study how different starting conditions or parameters of models affect the creative process. The final artifact produced by the simulation is also interesting. Computational simulation of creativity can be considered as a source of many artistic works [2][3].

Systems capable of enhancing human creativity using computational methods without necessarily being creative themselves are referred to as creativity support systems (CSS). These systems act as creative collaborators with scientists, designers, artists and engineers. CSS applies technology to assist humans to look farther and avoid thinking of the obvious concepts, that is, *thinking outside the box* and expanding the exploration boundaries.

There is widespread consensus that creativity involves two steps: a generative step to produce ideas; and a selective step to determine the most suitable ideas. Some CSS specialise on the generative step by producing a huge number of new ideas without selection criteria. The lack of assessment and selective ability has been one of the main criticisms of these systems. This is the case of the computational system for visual arts known as AARON [2]. Cohen reports that he can set the program before he goes to sleep, and wake up to find a hundred new and original images to review the next morning. The generative step is sometimes referred to as *divergent thinking*. Once a new idea emerges, the creator must determine its suitability. The selective step is referred to as *convergent thinking*. This stage draws on large amounts of domain knowledge to assess novelty and quality.

In the generative step, alternatives are normally generated by combining elements within and beyond the domain, but potentially good alternatives need to be chosen to avoid evaluating a vast number of possible combinations. This is achieved by reviewing existing processes in the domain and other processes belonging to other domains with subtle aspects in common. Creative people are skilled at finding these apparently different domains with common characteristics, and at pre-evaluating the alternatives (taking into account the relation between domains).

This paper introduces a formulation of creativity initially proposed in [5] and based on the central ideas of Boden's well-established theory on creativity [6]. In [5] the formulation is evaluated in cookery, and we argue that this formulation is powerful enough to be applied to other creative fields as music composition. The formulation enables us to focus on the selective step in creativity, evaluating alternatives from the relationship between appropriateness and relevance. We can even use concepts from different frameworks, that is, apply what Boden called transformational creativity.

The remainder of the paper is organized as follows: Section 2 provides a review of the literature on CC, and briefly lists previous works on CC applied to music. Section

3 introduces a novel transformational creativity approach based on the relationship between appropriateness and relevance concepts. The fourth section describes the experiment that enables us to understand how to implement the abstract concepts previously introduced in the music composition task. The final section gives conclusions and discusses future work.

2. Computational Creativity Approaches

One of the few attempts to address the problem of creative behavior and its relation with AI was done by Margaret Boden [6] [7]. She aimed to study creativity processes from a philosophical viewpoint focusing on understanding human creativity, rather than trying to create a creative machine.

Boden distinguishes between creativity that is novel merely to the agent that produces it, and creativity that is recognized as novel by society. The first is known as *P-creativity* (psychological creativity) and the second is known as *H-creativity* (historical creativity).

The most important contribution of Boden's study is the introduction of the idea of *conceptual space* that is composed of partial or complete concepts. She conceives the process of creativity as the location and identification of a concept in this conceptual space. The creative process can be performed by combining, exploring or transforming this conceptual space. According to Boden's theory, *combinational creativity* uses familiar ideas to generate a new idea in the form of unfamiliar juxtaposition; while *exploratory creativity* explores a conceptual space to create a new and unexpected idea. If the conceptual space is defined through a set of rules, when these rules change, then the process is called *transformational creativity*.

From Boden's study, it is not clear how the rules give rise to a particular conceptual space and, therefore, what is the true difference between exploring the space and transforming it. To clarify and to formalize the creative process, G. A. Wiggins [8] presented several papers emphasising the notion of search as the central mechanism for exploratory creativity and the notion of meta-level search related to transformational creativity. Wiggins posits a universe of possibilities \mathcal{U} which is a superset of the conceptual space. This universe is a multidimensional space, whose dimensions are capable of representing all possible concepts which are relevant to the domain in which we wish to be creative. For transformational creativity to be meaningful, all conceptual spaces are required to be subsets of \mathcal{U} .

Wiggins conceives exploratory creativity as a search of concepts in a specific conceptual space. The process involves three sets of rules that can be denoted as *acceptability*, *appropriateness* and *strategy*. The first set of rules is linked with belonging to the conceptual space. Moreover, acceptability is related to style. On the other hand, appropriateness rules are related to the value of the concept. Valuable concepts may become successful regardless of whether they are considered acceptable according to the acceptability rules. This second set of rules (which defines the value of a concept) is much harder to define because it depends on cultural and aesthetic aspects, specific context, personal mood, etc. However, it is important to note that, following Wiggins, appropriate means suitable to the task, but above all, original and surprising. Finally, there is a third set of rules linked to the search strategy. For instance, some people prefer to work 'top-down',

others 'bottom-up', while others rely on ad-hoc methodologies, using informed or uninformed heuristics and even randomness. Wiggins points out that separating acceptability and strategy rules can be used to describe situations where different designers, each with a personal way of finding new ideas, work within the same style (a shared notion of acceptability).

From Wiggins' perspective, the interaction of these three sets of rules (acceptability, appropriateness and strategy) leads to the exploratory creativity process. However, although working within three invariant sets of rules may produce interesting results, a higher form of creativity can result from changing these rules (transformational creativity). In other words, exploratory creativity consists of finding a concept in a specific conceptual space (following a specific strategy and assessing it by using a specific appropriateness set of rules), while transformational creativity involves the same process as exploratory creativity but changing the conceptual space, search strategy, or appropriateness assessment.

In addition to Wiggins' work, there have been other formalisations of specific aspects of the computational creative process [9][10] [11]. Although these formalisations are very helpful in clarifying the nature of creative computation and have given rise to some applications in domains such as graphic design, creative language, video game design and visual arts, the details of most of them are unspecified and the concepts they include are tricky to implement.

2.1. Computational creativity in music

The first creative computational systems were designed based on probabilities of note transitions and Markov-based techniques. The concept of probability transitions and Markov models was used to model musical styles by simply computing the note transition probabilities. Given a musical database corpus, new music can then be produced by generating notes using inferred probability distributions.

The ILLIAC (Hiller and Isaacson) (*Illinois Automatic computer*) was a series of computers that pioneered music composition based on Markov models. *Illiac Suite* is a 1957 composition for string quartet which is generally agreed to be the first score composed by a computer.

One of the most well-known applications of Markov chains for music generation is probably the Experiments in Musical Intelligence (EMI) designed by David Cope [13] although his musical results are not produced entirely automatically.

Not all the early work on composition relies on probabilistic approaches. Other approaches rely on simulating human composition processes using heuristic techniques. In [14] a more extensive list of previous works on computational creativity in music can be found.

3. A Novel Transformational Creativity Approach

A new formulation of transformational creativity was proposed in [5] that was based on the concepts of *framework*, *conceptual space*, *appropriateness* and *relevance*. This formulation considers a universal set of all concepts \mathcal{U} capable of containing concepts for every type of complete or incomplete artifact that might ever be imagined. A *framework*

\mathcal{F} is composed of a particular *H-conceptual space* $\mathcal{C} \in \mathcal{U}$ and two maps from \mathcal{U} to R , $a()$ and $r()$, called *appropriateness* and *relevance* maps respectively. Appropriateness is related to the success of considering a concept in this framework and relevance is a measure of the membership relation between the concept and the framework. In this formulation, originality and appropriateness are separated concepts. This separation is not clear in Wiggins formulation where appropriate means both suitable to the task, as well as original and surprising.

In this new formulation, we consider two classes of agents: *spectators* and *experts*. Any agent can act as a *spectator* and can easily obtain the relevance measure for any concept given any framework. However, only *experts* on a framework understand appropriateness for some concepts in the H-conceptual space.

Each expert i on a given framework \mathcal{F} is an agent that understands value $a()$ for concepts from some subset \mathcal{C}^i of \mathcal{C} . Inspired by Boden's theory, we call \mathcal{C}^i the *psychological* or *P-conceptual space*, that is, the concept space associated to the framework \mathcal{F} and to the expert i .

Experts only understand appropriateness for concepts from their P-conceptual space, and the values $a(x)$ for $x \notin \mathcal{C}^i$ are not understood by the expert. An expert can also have expertise in others frameworks. We use the notation \mathcal{C}_j as the H-conceptual space associated to the framework \mathcal{F}_j and \mathcal{C}_j^i as the P-conceptual space associated to framework \mathcal{F}_j and expert i . It is common that different frameworks share concepts; but obviously, appropriateness of the same concept can differ depending on the framework. Multi-expertise can be an advantage in the creative process.

We consider that given a framework, the appropriateness of a concept is independent of the expert. The difference between experts of the same framework is related to the different P-conceptual spaces (all being subsets of the H-conceptual space). In addition, we consider that both a framework and an expert can evolve -and both the H-conceptual space and the P-conceptual space can grow because of creative activity.

Contrary to appropriateness, relevance is the result of creative activities. Although evaluating the appropriateness requires some kind of talent or expertise, relevance evaluation can be easily performed by any agent (spectator) by means of an objective analysis of the framework. Thus a concept with high appropriateness in a framework is not necessarily highly relevant. In fact, an original concept always has low relevance in the considered framework.

The main problem regarding transformational creativity is the lack of knowledge from the expert's side on the appropriateness values for concepts outside his or her P-conceptual space. A CSS cannot directly obtain this appropriateness value. However, computational systems can be used for obtaining relevance values for any concept with respect to any framework, even if different to the framework task.

Given a set of different frameworks, $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_m$, and given a concept $x \in \mathcal{U}$, we consider the *relevance vector* of x with respect to the set of frameworks as $\phi(x) = (r_1(x), r_2(x), \dots, r_m(x))$, where $r_i(x)$ is the relevance of concept x with respect to the framework \mathcal{F}_i . This relevance vector contains indirect information regarding past creative activities involving this concept in those different frameworks. Our hypothesis is that no obvious relations between different frameworks exist, therefore the appropriateness of a concept x in a framework \mathcal{F}_0 , $a_0(x)$, and the relevance vector $\phi(x)$ are closely related. Concepts with similar relevance vector, $\phi(x)$, should have similar appropriate-

ness function. This hypothesis may not be true for a small set of frameworks but, from our previous experiments, it seems to be true for larger sets.

4. Experiment description

To illustrate the implementation of the approach presented in section 3, let us consider the task of composing a single voice tune of a specific style, for example a *reel* (a folk rhythm originated in Scotland).

We have considered single voice folk songs written using ABC notation and our goal is to substitute part of one song by other compatible sequence of notes with similar relevance vectors.

We start with an already existing reel and we extract a small part (a sequence of notes) of this reel. We then substitute this sequence of notes by another compatible sequence (that follows some structural rules of harmony and style) and which has a similar relevance vector as the original part. To this end, we considered that the position of the extracted sequence, and the type of tune to be composed (reel, in this example) is the framework. To obtain the relevance vector, we consider other frameworks (other styles of tunes such as 'jig', 'waltz' or 'polka'). The hypothesis is that substituting the musical element (concept) by another with similar relevance vector, the appropriateness measure will be similar.

4.1. ABC notation

ABC notation is a text-based music notation system popular for transcribing, publishing, and sharing music, particularly on the internet. It was formalised (and named) by Chris Walshaw in the early 90s with help and input from others.² Since then, ABC has gone from strength to strength and is widely used by folk musicians, especially from Western European origin, e.g. English, Irish, Scottish, and which typically produce single-voice melodies. The most recent standard for ABC is v2.1, released in 2011.³

Each ABC tune consists on a tune header and a tune body (it may also contain comment lines or stylesheet directives). The tune header is composed of several information field lines containing the reference number, title, tempo, default note length, rhythm, key, etc. The tune body contains the notes, duration, bar lines and other musical symbols.

4.2. The notes in ABC notation

Starting at middle C, the notes in that octave are shown as CDEFGAB. The next octave is shown in lowercase and the next one using apostrophes. The octave immediately below middle C is shown by a comma immediately following the note name (see Figure 1).

The range can be extended further by adding more commas or apostrophes. Sharps (\sharp) and flats (\flat) are shown using circumflex (^) and underscore (_) before the note name. Different key signatures enables sharpening or flattening all notes of the same name. The equal sign (=) is used to naturalise (\natural) a note.

²<http://abcnotation.com/contact>

³<http://abcnotation.com/wiki/abc:standard:v2.1>




Figure 1. A scale in standard notation, with the equivalent abc note names shown.

ABC enable us to set the 'default note length' for each tune. If the particular note length is a divisor of the default length, then it is shown as '/n' following the note name. If it is a multiple of the default length it is shown as 'n' following the note name. The symbols '>' and '<' are a shorthand notation for dotting. They transfer half the value of the note on one side to the other side.

In the tune body, spaces and bars are used to facilitate the reading, double quotes indicate some information such as accompaniment chords and other symbols such as tilde or curly brackets indicate the presence of an ornament.

In this experiment we only have considered single voice tuned without accompaniment chords and we have dispensed with all note information from a tune except pitch and length. This means that, for instance, from this ABC code:

```
K : D
| : "G"ABC2|A > AA2|GAdc : |
```



we extract the following information:

$$\begin{pmatrix} A & B & c\# & | & A & A & A & | & G & A & d & c\# \\ 1 & 1 & 2 & | & 1.5 & 0.5 & 2 & | & 1 & 1 & 1 & 1 \end{pmatrix}$$

where the first row of the matrix corresponds to the pitch notes and the second row to the respective lengths. Vertical lines separate sequences of four default times (default=eighth note)

4.3. Database description

The database is formed by 2386 tunes classifying in 16 different rhythms: reel: 1002, jig: 484, polka: 164, song: 152, hornpipe: 138, slide: 98, slip jig: 76, waltz: 65, march: 35, carolan: 31, slow air: 29, highland: 24, strathspey: 21, hop jig: 20, set dance: 20).

This database is extracted from databases:

<http://www.folktunefinder.com/tunes>

<http://abcnotation.com/>

<http://www.norbeck.nu/abc/>

all of the tunes are single-voice melodies in different key signatures.

4.4. Procedure description

The *concepts* in our experiments are composed by two sequences: a sequence of pitches and a sequence of note lengths that add up to four. The idea is to substitute this sequence of notes by another with similar relevance vector. By repeating this step with different sequences, we can obtain a completely different song starting from a original song.

As we have 16 different types of rhythms, we consider relevance vector of dimension 16. Each component of the relevance vector for a sequence of pitches and lengths is calculated by counting the percentage of sequence occurrences in the collection of tunes of a specific rhythm. In this calculus, we give different value weights if all pitches and lengths occur, or only the sequence of pitches or the sequence of note lengths appear. Although it is a parameter of the method, in the preliminary experiments we have obtained good results using the following scores:

$$\begin{pmatrix} A & A & A \\ 1.5 & 0.5 & 2 \end{pmatrix} \rightarrow 100 \text{ points}$$

$$\begin{pmatrix} A & A & A \\ \cdot & \cdot & \cdot \end{pmatrix} \rightarrow 20 \text{ points}$$

$$\begin{pmatrix} \cdot & \cdot & \cdot \\ 1.5 & 0.5 & 2 \end{pmatrix} \rightarrow 10 \text{ points}$$

If the sequence of pitches and lengths is found in a specific tune of a certain style, we add 100 points to the respective relevance coordinate, if the sequence of pitches is found but with different note length, only 20 points is added. Finally, if the sequence of note length is found with different pitches, only 10 points is added. Only one of the four scores (100, 20, 10 or 0) is considered for each song, and the sum of all scores for all songs belonging to the same style is normalized by the number of songs of this style. This calculus performed for all styles produces the relevance vector of this sequence of pitches and lengths.

The next step is to consider different sequences of pitches and lengths to replace the previous one. We consider sequences of different pitches and lengths (almost one hundred of them) that satisfy some harmony constraints (related to the key) and rhythmic constraint (related to the style). At the moment, this step is performed manually. For each of the sequences, the relevance vector is obtained using exactly the same algorithm explained in the previous paragraph.

After all relevance vectors are obtained, the Euclidean distance between each and the relevance vector of the original sequence of pitches and lengths is computed. The generated sequences are sorted by distances, and those with a shorter distance can be selected to replace the original one. The more sequences replaced, the more the resulting tune will differ from the original.

5. Conclusion and future work

The previous section explains in detail the implementation of the transformational creativity approach. Our goal is not to create a computational system capable of composing a creative piece ready to be interpreted in a concert, or to be included in an album for sale. Conversely, we want to test the relation between the concepts of appropriateness and relevance in this field. To better understand the scope of our formalization we have chosen a simple musical environment and eliminating complex ornaments.

It is difficult to measure the quality of the melodies obtained. We are working on creating a website to publish some of these melodies and also a web application that can change some of the parameters such as the sequence length replaced, score system, the harmonic and rhythmic constraints).

Finally, we are working on automating the generated sequences step since it is the most onerous part of the whole process.

6. Acknowledgments

This research has been supported by the INVITE Research Project(TIN2016-80049-C2-1-R and TIN2016-80049-C2-2-R (AEI/FEDER, UE)), funded by the Spanish Ministry of Science and Information Technology.

References

- [1] C. W. Taylor. Various approaches to the definition of creativity, in R. J. Sternberg (ed.)*The Nature of Creativity*, Cambridge University Press, 99125, 1988.
- [2] H. Cohen. *The Further Exploits of Aaron, Painter*. Stanford Hum Rev. vol 4, n. 2 pp 141-158. 1995.
- [3] J. D. Fernandez, F. J. Vico. *AI Methods in Algorithmic Composition: A Comprehensive Survey*. Journal of Artificial Intelligence Research. 48: 513582.
- [4] D. Cope. *Computer Models of Musical Creativity* (MIT Press). December, 2005.
- [5] F. Ruiz, C. Raya, A. Samà, N. Agell. Transformational creativity tool to support chocolate designers Pattern recognition letters Vol. 67, num. Part 1, p. 75-80. 2015
- [6] M.A. Boden, *The Creative Mind: Myths and Mechanism*. Weidenfiel and Nicholson, London, (1990).
- [7] M.A. Boden, What is creativity? In M.A.Boden, editor, *Dimensions of Creativity*. MIT Press, 75-118. 1996
- [8] G. Wiggins, A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19 pp. 449-458. 2006
- [9] G. Ritchie, Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17 (1), 67-69 (2007).
- [10] G. Ritchie, A closer look at creativity as search, In *Proceedings of the 3rd International Conference on Computational Creativity* (2012).
- [11] J. Charnley, A. Pease and S. Colton. On the Notion of Framing in Computational Creativity, In *Proceedings of the 3rd International Conference on Computational Creativity* (2012).
- [12] R. Manurung, G. Ritchie and H. Thompson. Using a genetic algorithms to create meaningful poetic text. *Journal of Experimental & Theoretical Artificial Intelligence* 24(1), pp43-64. (2012),
- [13] D. Cope. *Experiments in musical intelligence*. Madison, WI: A-R Editions. (1996).
- [14] R. López de Mántaras. "Computational Creativity". *Arbor*, 189 (764). (2013)