Darwinci: Creating Bridges to Creativity

Francisco C. Pereira[1], Penousal Machado[2], Amílcar Cardoso[1]

[1]Departamento de Engenharia Informática da Universidade de Coimbra
{camara, amilcar}@dei.uc.pt
[2]Instituto Superior de Engenharia de Coimbra
machado@dei.uc.pt

Abstract. This paper presents Darwinci, a system that generates new ideas, using a multi-domain knowledge base composed by musical and drawing structures. Its theoretical background comes from the theory of Divergent Production as stated in [7], Genetic Algorithms and metaphor interpretation. Rather than attempting to model human thought, Darwinci's main goal is to make cross-domain transfer of ideas, in order to create something new, thus becoming a good proposal for creative problem solving.

1 Introduction

It is known and stated that some of our most brilliant ideas come up when our mind "wonders" around something that has nothing to do with the problem we are trying to solve. Due to the complexity of the world and to its representation in computers, the elaboration of a "divagation" machine has never appeared, as far as we know, in AI literature. However, several psychologists agree that this capacity of "wondering" is very important in solving problems that involve some degree of creativity [3].

The work presented here is part of an ongoing research project, centered on an important theory of creativity: J.P. Guilford's theory of "divergent production". More specifically, this paper presents an implementation of what Guilford calls "the ability to generate variety and amount of information". Darwinci uses a genetic algorithm (GA) to produce new ideas, supported by a multi-domain knowledge base. Each domain will be composed of two types of information: a set of tree-like structures representing "individuals" (e.g. musical pieces, drawings), a set of conceptual nets to assert knowledge about it. This representation can be used with a variety of domains, including Music and Visual Arts (our experimental domains). As will be shown, we find several conceptual and structural connections among these forms of art.

The paper is organised as follows: Section 2 is dedicated to a small survey and dissertation on the study of creativity. We believe that the importance of this area touches not only psychology but also AI. In sections 3 and 4, we will describe our previous work in Music and Visual Arts. The inter-crossing between the two domains will be addressed in section 5. Finally, in section 6, we draw some generic considerations about Darwinci, and point to some unexplored aspects in this field.

2 Studying Creativity

The study of creativity has motivated research from various viewpoints like the creative product; the creative process; the creative environment; the creative person [3]. Psychologists like [7], [5] have dived into the search for answers to questions like: How do we create novel ideas? Why are some people more creative than others? Are there any special traits directly related to this capacity? How to measure it?

In AI many relevant examples have already emerged, be it related to analogy [6], metaphor and linguistics [15], [6], neural networks [8], or creativity modeling and studying [2]. Whether in search for new problem-solvers, for ways of understanding cognition or programs that create interesting things, the motivations behind these studies touch this ability of ours: creativity.

Without making claims on how do we create or what is a creative product, our interest is the computational application of theories of creativity in order to create new ideas.

The work we present in this paper is connected to some ideas of Guilford [7]. In his theory - Structure of Intellect -, divergent production is considered as the ability to "generate variety and amount of information; most involved in creative potential"[3]. More specifically, Michael [11] states that this can be achieved by recurring to the process of "transform recall", in which the same information is applied to different contexts (as if one is trying to apply an old puzzle piece into a new one). According to Guilford, creativity is directly related to our capacity to interconnect and correlate apparently divergent ideas, i.e., to see similarity where, by default, there seems to be none (e.g. the famous example of the self-biting snake and the benzene molecule of Kekulé's dream). Computationally, this can only be achieved with a high degree of flexibility in the representation and some tendency to divagate through unrelated areas, which means that we must enable associations between concepts from various areas. A program able to do it must accept an open representation (in the sense that it must allow different domains) and have a mechanism to operate on the knowledge itself without loosing semantics.

Our approach to computational production of (as possible creative) ideas is mainly based on: divergent production and operations on knowledge.

We have chosen two domains in which we have some experience and where creativity is relatively unconstrained (in opposition to Natural Language, for example). These are Music and Visual Arts. We have already developed systems under these two subjects: SICOM [12] and NevAr [10].

In both works, structure and hierarchy play an important role and we believe that, at least in upper levels of abstraction, one can surely try this cross-domain divergent production between Music and Visual Arts. Parallels between these two forms of expression are neither few nor new. It's common sense that there is strong intersection between the several forms of art. In fact, history tells us that a given style is not associated to a single form of art, its general abstract characteristics are spread through various artistic and geographic areas. Maybe due to this, specific concepts of each form of art are often applied to others, yielding new and fertile metaphors. Concepts like "colour", "texture" and "contour" are often present in Music. Likewise, “rhythm”, “dissonance”, are used in Visual Arts. Other concepts like harmony, motive, tension or dynamics are equally used in both domains. As we do not want to play paintings or to paint musical pieces, we will not “force” correspondence between those concepts.

As will be seen bellow, the connections between domains are generated through a mechanism derived from the work on metaphor of [15] and [6]. In the next two sections, we will describe previous work on Music and Visual Arts, switching gradually to the main scope of the paper.

3 Creating Music

The analysis and composition of Music by computers has already some history. Its state-of-the-art, by now, comprehends several works, both in analysis [14] and in composition [4].

Following ideas from [4] and [9], our method of generation of music is based on the reusability of previous ideas. Although in former work [12] we were interested in creating entire pieces, by now and for the present purpose, we are focussing on the creation of chord sequences with rhythm. A piece of music can be represented by a hierarchical structure corresponding to its harmonic, melodic and rhythmic analysis. Creating a new musical piece, consequently, can be seen as constructing such an organisation. So, these chords sequences come out from the construction of a tree-like structure, much in the same shape as in [9] where, analytically, one can represent an entire music with a hierarchical structure defined by grouping, time-span and metric rules. In the next paragraphs, we will describe briefly some aspects of our musical system (SICOM), which generated new pieces through the use of Case-Based Reasoning [12].

SICOM used pre-elaborated analysis of music coded as trees, with non-hierarchical links between nodes, used for explaining relations among them (see fig.1).

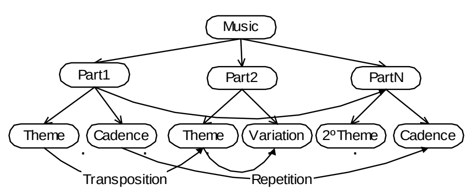

Fig 1. A Musical Structure

In the act of producing new structures, the system used these links as “suggestions” with associated strength (for example, Repetition may be strong and Transposition may be weak, in fig.1) to reduce the search space and to keep some coherency throughout the piece.

To make a synthesis of the SICOM’s process of generation, the algorithm takes four steps: (1) Search for an acceptable node in the memory; (2) Apply it and, if needed, adapt it; (3) Spread its “suggestions”; (4) If no more nodes are expected, finish, else go to 1.

The last condition is simply accomplished by structural links, i.e., when choosing a node, its structural connections to descendants are also spread, becoming strong suggestions. If there are no more of these, the piece is finished. Since explaining SICOM in detail is not the scope of this paper, we redirect the reader to [12] for further details. For present purposes, we retain some of this system’s characteristics, namely:

- Representation – Although applying some changes in order to approach closer to Lerdhal and Jackendoff’s theory enhancing concepts like “tension” and “relax”, the basics of the former representation are kept.
- Domain adaptation – While developing SICOM, we have built some adaptation algorithms. Although highly biased by the musical idioms we worked with, we keep the main ideas.
- Some theoretical background – We have already applied some ideas from Guilford [7] in SICOM. Then, it was centred mainly on work on the similarity metric. Presently, we are following different paths to explore these ideas.

4 Creating Images

The basic idea for our approach to the generation of images is based on an ongoing research project [10]. The idea is using a genetic algorithm to generate aesthetically pleasing images, which has already been shown to be effective [13][1]. The main problem of this approach lies on the assignment of fitness values to the individuals. Generally, the fitness is supplied by the user (e.g. [13]). In [10][1], the evaluation of the individuals is made through artificial neural networks.

Our approach to the “evaluation problem” is radically different, and will be described in the next section. For the scope of this section, let’s consider that evaluation is made by a black box. Apart from this, our algorithm is very close to the one presented in [10]. The differences are related to representation, resulting in minor changes to the crossover and mutation algorithms.

4.1 Overview of the algorithm

One of the innovative aspects of the algorithm presented in [10] is the use of background knowledge. This is achieved in two different ways:

1. The initial population does not need to be random. Any set of images, including famous artworks, can, potentially, be used as initial population.
2. It maintains a knowledge base (KB) of images, which are not subject of selection. The images in this KB can be selected for mating with others from the current population or between themselves. The KB can be initialized with any set of images we choose. Additionally, individuals generated by the system can be added to the KB.

The use of a KB has additional benefits, namely, the increase of population diversity and preventing the early lost of remarkable individuals [10].

Let's make a brief overview of the generation algorithm: we start with a working set of images built by the KB and the current population. The individuals in this set are already evaluated.

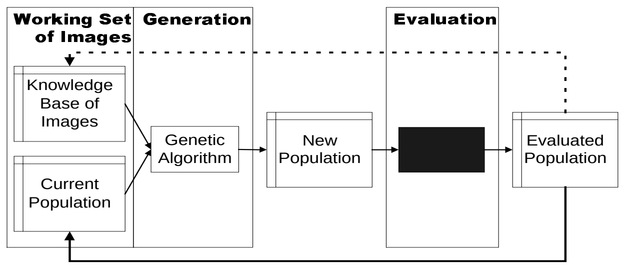

Fig. 3. The generation algorithm

From the working set, the GA creates a new population through genetic recombination of the selected individuals (probabilistic selection based on fitness). The images of this new population are evaluated, resulting in an evaluated population, which becomes the current population. Additionally, images that have an aesthetic value higher than *c*, or superior to the average of the population by a value *d*, are added to the KB. A new working set is achieved and the process is repeated.

4.2 Image representation

The representation issue is one of the most important ones in this type of system. One common approach [13][10][1] is to represent the images through mathematical functions. This representation has several drawbacks. We are limited by the representational power of the used formulas, and there is no procedure for converting a generic image to this type of representation. This creates a great problem to us since we want to use artworks made by human artists. A further problem results from the fact that we don't have any explicit representation of the basic building blocks of the image nor of its relations. Considering these problems, we chose a different representation, which is quite similar to the one used in the music domain.

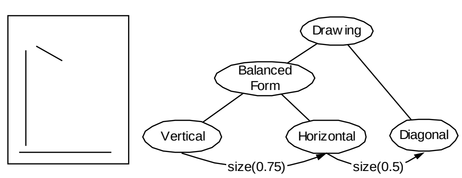

Fig. 4. A simple drawing and its representation.

This representation is still tree-like, allowing us to use the same type of crossover and mutation operators. In our previous work, the crossover of two individuals consisted in the exchange of sub-trees between these individuals. Our new crossover operator also exchanges sub-trees. However, since we have relationship links, these can be broken during crossover. When this happens, the algorithm searches for new relations so that the broken links can be reestablished. The same happens during mutation.

5 Mixing Structures

The organisation of Darwinci corresponds to the following diagram:

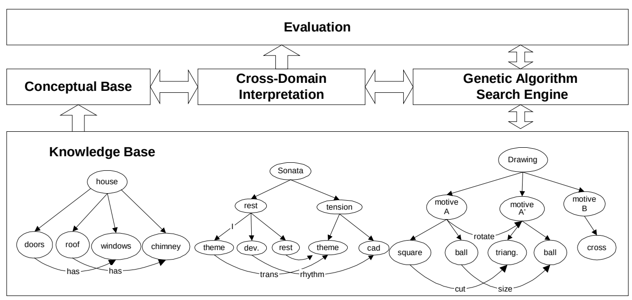

Fig. 5. General architecture of Darwinci

The basic blocks of Darwinci can be described as follows:

Knowledge Base – Set of structures from a variety of domains. Its representation is similar to the one described in previous sections.

Conceptual Base – Composed by the conceptual graphs describing domain knowledge, organized according to an implicit ontology (easily extractable by following “isa” arcs). These graphs allow inter-relations among any pair of concepts.

Genetic Algorithm Search Engine – This algorithm is similar to the one described in the previous section. The only difference lies in the fact that we are dealing with individuals from different domains. This, however, doesn’t imply radical changes to the algorithm, since the representation is kept among domains. Therefore, the main difference lies in the fitness function. In Darwinci, the Evaluation module determines the fitness values. During mutation and crossover, some of the previously existing relation links may be broken. The algorithm tries to re-establish links through the search of new relations, which must be coherent with the conceptual. This search is made in a conservative way; i.e. it tries to maintain the previous relations with the minimum amount of change. This is made through the use of the Cross-Domain Interpretation Module.

Cross-Domain Interpretation – This is one of the key blocks of Darwinci. When interpreting a structure, this module establishes cross-domain links between the domains in question; i.e. it establishes a metaphor between the domains [15] and [6]. The individual will be interpreted according to this metaphor (e.g. “colour gradient is harmonic progression”). It creates correspondences that control the translation of the structure to a specific domain. The nodes of a given structure are all interpreted in accordance to the same metaphor. In other words, when using a metaphor (like “harmony is color contrast”), the system is supposed to create new connections, (like “chord is color”, “Dominant Major is Blue”) and use them to translate nodes.

Evaluation – This module is responsible for the assignment of fitness, which relies on the coherency of the individual with the conceptual nets present in the system. We can test coherency with respect to several domains (e.g. music, painting) or we can “direct” the search to a specific domain. To assess coherency, the Evaluation Module calls the Cross-Domain Interpretation. Following the current metaphor links, it finds concepts that can evaluate coherence (suppose, in music, the “Sonata form” net, which would have a set of concepts that characterise it: “has_many cadences”, “Middle_Section is Dominant”, etc.) and uses it to calculate the fitness.

6 An Example

In this section, we show a small example of Darwinci. Suppose the conceptual base corresponds to fig. 6 and that we have an initial knowledge base as in fig. 7.

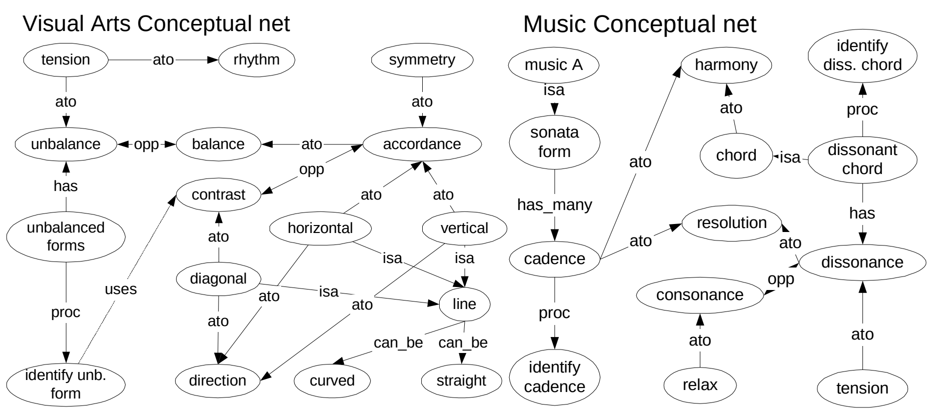

Fig. 6. Musical and Visual Arts Conceptual nets. “ato”–associated to; “ako”–a kind of; “isa”–is a; ”proc”-procedure; “opp”–Opposite

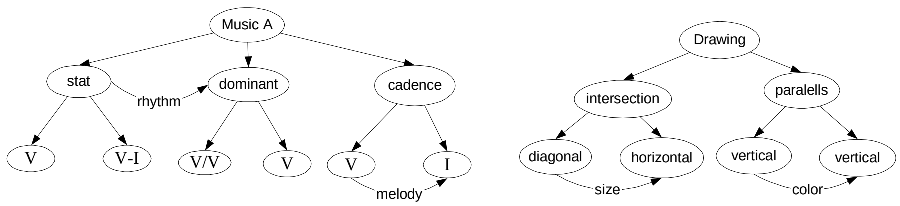

Fig. 7. Two structures of Music and Drawings

The GA starts by selecting a random set of individuals from the knowledge base. Let's assume it picked the two structures from fig. 7, and that the following individuals were created through recombination:

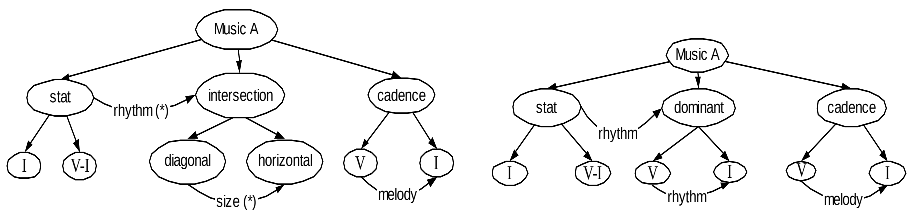

Fig. 8. On the left, the resulting structure, before inter-domain "translation". (*) Broken link. On the right, the final structure

In the following step, the cross-domain interpretation, Darwinci establishes new "metaphorical" links between conceptual nets. (fig. 9).

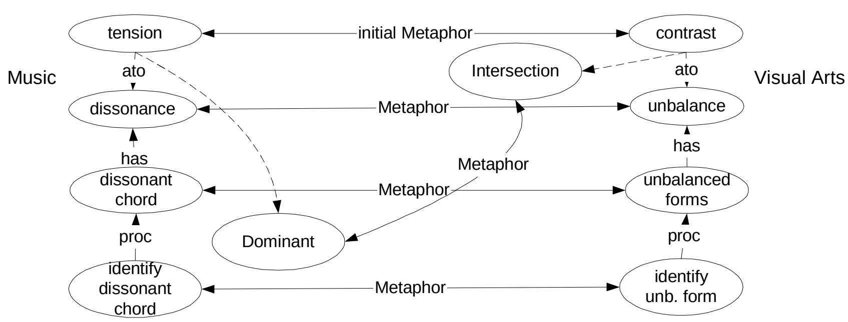

Fig 9. The metaphor links have been spread across both nets

As can be verified, Darwinci has (in this example) an initial metaphor between "tension" in music and "contrast" in visual arts. Departing from this, it creates a structure mapping that will be used as the base for interpretation. Following the systematicity principle [15], some of the remaining "translations" are generated (assuming a generated metaphor link between "main degrees" and "direction", the result could be "Horizontal = I", "Vertical = IV" and "Diagonal = V").

Still in the cross-domain interpretation phase, the program tries to re-establish links. One of the ways to do this is simply by reapplying it to the new situation. Suppose this is what happened to the first broken link (rhythm). The other way to re-establish links is by following metaphoric links. Let's believe it would follow them and detect a link between "rhythm" and "size". Then it would be recovered with success. The new structure is presented in Fig. 8 (right).

The next phase is the evaluation. This critical step is also supported by the metaphor links. As can be seen, the concept "has_many cadence" is associated to "Sonata", which is associated to "Music A". According to this, Darwinci executes the procedure associated to "cadence", counting the number it "succeeds". This, with other concepts (like "dominant in middle_section") will help the program to evaluate the fitness.

Following this algorithm, the program is expected to arrive at some interesting solutions. Some of these steps are not yet implemented, but the basic ideas were shown. For the sake of simplicity, this example is highly biased (there can be many unfortunate interpretations and crossover results) but, with the help of the GA, Darwinci is expected to get solutions like the one above sooner or later.

7 Conclusions and Further work

This paper presented the overall architecture and theoretical background behind a work on computational creativity. As in other works, it is natural that implementations and further investigations mutate the original idea gradually until getting the final result. But we think this is a rather mature framework at least in what concerns to the idea of using a GA to create a multiplicity of structures and using "metaphorical" connections between conceptual webs to make cross-domain interpretation. The authors have already achieved some confidence on working on the generation of musical and drawing structures in previously published work.

Some parts of this work, namely the conceptual nets, need an extensive search in order to achieve an acceptable degree of coherency and completeness. In order to do it, and to allow relatively low complexity, we must assume that there must be a limited set of relations among concepts. This will introduce some bias in the cross-domain interpretation, but we hope to keep some expressiveness on these links.

Presently, the cross-domain interpretation and the conceptual base modules are being implemented. We are following previous work by [15] and [6]. After the complete interpretation of this system, we envisage the following two further steps: the application to other domains (like story plots and dance) and the design of a learning module (at the level of the conceptual base module).

Without aiming to obtain results of excellence in the creation of artworks, we hope to shed some light on the search for creative production in Artificial Intelligence.

Acknowledgements

We would like to thank CISUC (Centro de Informática e Sistemas da Universidade de Combra) for supporting our research. The first author is sponsored by a grant PRAXIS XXI BD 9611/96.

References

1. Baluja, S.; Pomerlau, D.; Todd, J.; 1994, *Towards Automated Artificial Evolution for Computer-Generated Images*. Connection Science, 6(2), 1994, pp.325-354.
2. Boden, Margaret; 1990. *The Creative Mind: Myths & Mechanisms*. Basic Books.
3. Brown, R. T. ; 1989. *Creativity: what to are we to measure?* In Glover, J.; Ronning, R. and Reynolds, C., Handbook of Creativity, Plenum Press, New York.
4. Cope, D.; 1996; *Experiments in Musical Intelligence*. Vol.12 of the collection: The Computer Music and Digital Audio Series. A-R Editions.
5. De Bono, E.; 1986. *El pensamiento lateral*. Manual de la Creatividad. Barcelona, Espanha: Padóis.
6. Gentner, D.;1983. *Structure Mapping: A Theoretical Framework for Analogy*. Cognitive Science, 7(2), pp 155-170.
7. Guilford, J.P. 1967. *The Nature of Human Intelligence*. New York: McGraw-Hill.
8. Holyoak, K. J. and P. Thagard. (1989). *Analogical Mapping by Constraint Satisfaction*, Cognitive Science 13, pp 295-355.
9. Lerdhal, F. and Jackendoff, R. 1983. *A Generative Theory of Tonal Music*. Cambridge, Mass.: MIT Press.
10. Machado, P.; Cardoso, A.; 1997. *Model Proposal for a Constructed Artist*. In Proceedings of World Multiconference on Systemics, Cybernetics and Informatics, SCI'97. Caracas.
11. Michael, W.B.; 1977. *Cognitive and affective components of creativity in mathematics and the physical sciences*. In J.C. Stanley, W.C. George & C.H. Solano(Eds.), The gifted and the creative: A fifty-year perspective. Baltimore: Johns Hopkins University Press.
12. Pereira, F.; Grilo, C; Macedo, L; Cardoso, A.; 1997. Composing Music with CBR. In Proceedings of the First International Conference on Computational Models of Creative Cognition. Dublin.
13. Sims, K.; 1991. *Artificial Evolution for Computer Graphics*. SIGGRAPH'91, Computer Graphics, Vol. 25, Nr.4, pp. 319-328, ACM.
14. Smaill, A.; Wiggins, G.; and Harris, M. 1993. *Hierarchical Music Representation*. In Computers and the Humanities, vol. 27.
15. Veale, T. ; Keane, M. T.; 1994. *Metaphor and Memory: Symbolic and Connectionist Issues in Metaphor Comprehension*, in the Proc. of the European Conference on Artificial Intelligence Workshop on Neural and Symbolic Integration, Amsterdam