德古意特认知语言学应用丛书

CREATIVITY AND ARTIFICIAL INTELLIGENCE

A CONCEPTUAL BLENDING APPROACH

创造力与人工智能的概念整合研究

Francisco Câmara Pereira 著

上海外语教育出版社
外教社 SHANGHAI FOREIGN LANGUAGE EDUCATION PRESS
www.sflep.com

德古意特认知语言学应用丛书

CREATIVITY AND ARTIFICIAL INTELLIGENCE

A CONCEPTUAL BLENDING APPROACH

创造力与人工智能的概念整合研究

Francisco Câmara Pereira 著

上海外语教育出版社
外教社 SHANGHAI FOREIGN LANGUAGE EDUCATION PRESS
www.sflep.com
DE G De Gruyter Mouton

图书在版编目（CIP）数据

创造力与人工智能的概念整合研究 /(丹)弗兰西斯科 · 卡马拉 · 佩雷拉著.
—上海：上海外语教育出版社，2017
（德古意特认知语言学应用丛书）
ISBN 978-7-5446-5011-3

I. ①创… II. ①弗… III. ①认知语言学—应用—创造能力—研究—英文
②认知语言学—应用—人工智能—研究—英文 IV. ①H0-06

中国版本图书馆CIP数据核字（2017）第231545号

图字：09–2017–666

出版发行：上海外语教育出版社
（上海外国语大学内） 邮编：200083
电　　话：021-65425300（总机）
电子邮箱：bookinfo@sflep.com.cn
网　　址：http://www.sflep.com
责任编辑：许进兴

印　　刷：上海叶大印务发展有限公司
开　　本：700×1000　1/16　印张 16.75　字数 317千字
版　　次：2018 年 5月第 1版　2018 年 5月第 1次印刷
印　　数：1 500 册

书　　号：ISBN 978-7-5446-5011-3 / T
定　　价：55.00 元
本版图书如有印装质量问题，可向本社调换
质量服务热线：4008-213-263　电子邮箱：editorial@sflep.com

出版说明

认知语言学是语言学的一门重要分支学科，自20世纪80年代诞生以来，受到了国际和国内学界的广泛关注。近年来，外教社陆续推出了一系列相关丛书，集中体现了国际、国内的优质研究成果。其中“国际认知语言学经典论丛”收入了Ronald Langacker、Leonard Talmy、Dirk Geeraerts等国际认知语言学领域顶尖学者的经典作品；“外教社认知语言学丛书·普及系列”、“外教社认知语言学丛书·应用系列”则体现了国内学界的最新研究成果。这些丛书因内容权威、见解独到受到了外语界的广泛好评。

认知语言学作为一门新兴的跨领域学科，与多种学科有密切的联系，具有很强的应用意义。因此，我们又从德古意特出版社近年来推出的相关学术图书中精选了6种，组成“德古意特认知语言学应用丛书”，引进出版。丛书反映了近十几年来认知语言学应用领域的前沿成果，集中体现了该学科的理论与实践、应用与展望，及其与人工智能、诗学、语言教学等领域的联系和互动，信息量大，时效性强，代表了国际认知语言学应用研究方面的最高水准。丛书作者汇集了Gitte Kristiansen、Francisco J. Ruiz de Mendoza Ibáñez、René Dirven等国际认知语言学界领军人物，以及欧美相关领域的优秀学者，体现了国际认知语言学应用研究方面的最强阵容。

相信丛书的引进可进一步帮助国内读者了解这一领域的研究热点和最新成果，为国内研究者带来新的启示。

Creativity and Artificial Intelligence

A Conceptual Blending Approach

by
Francisco Câmara Pereira

Acknowledgements

In parallel to the work behind this book, a whole novel could have been written, with many more actors than appears in the final document. The truth is that, without these actors, it would certainly lose much of its value.

My former supervisor and colleague, Prof. Dr. Amílcar Cardoso, has been a key actor from the first to the last chapter, sometimes as a patient advisor and listener, other times as a friend.

Some colleagues (and friends) had the leading role on occasion, even if they were unaware of it. In fact, a substantial part of this book was a result of talks, collaborations and constructive arguments with Tony Veale, Pablo Gervás, João Leite and Geraint Wiggins.

Also, without the rest of the AILab crew (before and after ECOS), a lot of science and friendship chapters would certainly be irreparably lost. These were Penousal, Paulo, Grilo, Xico, Leonor, Prof. Carlos Bento, Prof. Ernesto Costa, Jorge Tavares, André, Miguel Ferrand and Ana Alves.

Special thanks must be made to Paulo Ribeiro, João Pedro Martins, Bruno Leitão and Bruno Marques who have contributed with work that is also strongly linked to the thesis that originated the book.

A large part of this novel was done within the environs of the Departament de Informatics Engineering of the University of Coimbra, and I must not forget to thank my colleagues there and particularly the staff, Isabel, Paula, Manuela and Márcia, for their support.

The other side of the story, with different scenery and people, starts at home. Ana has given the strongest support and love as only she can give and has frequently drawn from a limitless reservoir of patience. Kali, Becas and Sansao have also been patient and supportive, in the way that only cats and dogs can be. Kali and Becas have left us during this period, and we feel a deep sorrow, for this and future novels would never be the same without them.

Finally, my parents and sister deserve a special mention for the many times that I missed them, especially when they needed me. And for the support they have given from the beginning.

I must also kindly thank the reviewers of this (and previous) publication(s) for the several constructive and incisive comments that definitely helped improve its potential as a contribution to science.

Contents

Chapter 1
Introduction

> It is obvious that invention or discovery, be it in mathematics or anywhere else, takes place by combining ideas... (Hadamard)

> Pour inventer il faut penser à côté. (Souriau)

> The useful combinations [of ideas] are precisely the most beautiful. (Poincaré)

1. Motivations

Right from the start, the main focus in AI research has always been with the issue of problem solving. Seen from this point of view, intelligence corresponds to the ability to solve (complex) problems, from the accurate autonomous movement of a robot arm to the understanding of a natural language sentence. The *classical* setting is that of a search in a *space of solutions* for the problem, where an *intelligent agent* looks for the best choices. A commonly used analogy is of this agent travelling in a *search space* with mountains (highly valued solutions), valleys (bad solutions), islands (mountains surrounded by valleys) and plains (areas where quality of solution hardly changes)[1].

One of the most common criticisms made of Artificial Intelligence methods of problem solving is their limited ability to deal with situations not predicted in the specification. The search space is normally strictly defined, however flexible, complex and adaptable the system seems to be. When facing a problem with no satisfactory solution in its search space, an AI system simply returns, at best, the least unsuccessful result that exists in that search space, even when the solution is achievable via the simplest operation of changing perspective, relaxing a constraint or adding a new symbol. In other words, such systems are hardly capable of performing what we normally call *creative behavior*, a fundamental aspect of intelligence.

However, the recognition that there is a deficit of creativity within AI systems does not by itself bring new solutions any more than it reasserts that computers, as we know them, are formal machines that are limited to their closed worlds. The question therefore arises about what can be done to make them more creative or even if, with the current computational architectures, that is possible at all. To some extent, some of the current state-of-the-art paradigms (such as Evolutionary Computation, Multi-agent Systems or Case-Based Reasoning) have been responsible for many of the developments regarding the first part of the question. Indeed, we have been developing less rigid systems in previous years and sometimes even producing striking results. Nevertheless, when any of these systems finds a situation for which it was *a priori* not specified to solve, it is definitely not able to cope with it.

The second half of the question concerns primarily what the essential components of a creativity model could be and whether these can be present in a formal machine. There is no definitive answer for this, yet we can allow ourselves cross-fertilization from other areas, such as Psychology, Cognitive Linguistics, Cognitive Science and Philosophy, in the speculation and building of a possible solution.

The relationship between Intelligence and Creativity poses further questions. Are these two independent properties of cognition or, on the contrary, are they interrelated and inseparable? More specifically, if taking a traditional AI perspective: isn't creativity about search? Is it a different approach to intelligence?

These questions are present throughout this book, which is an attempt to approach them according to a perspective that, while centered on Computer Science and AI, also lifts contributions from other areas.

The reader can find most of the material described and used in the context of this book in the attached CD and at `http://eden.dei.uc.pt/~camara/AICreativity`. In this website, the reader can find the latest additions to this and related works.

2. Overview

There is general agreement that the ability to find relations between apparently unrelated knowledge is a creative activity. As can be found in many studies from the area of cognitive psychology, the creative faculties of the human mind are hugely dependent on the ability to search through spaces or "viewpoints" that are different from the ones immediately or more obviously involved. For example, according to (Marin and De La Torre 1991), our capacities of abstraction, symbolic analysis, of finding less-obvious relations, among others, are associated to creative production. Indeed, many important discoveries, great music compositions or paintings were reportedly achieved at times of wandering in domains not directly related to the actual problem (e.g. the famous dream of Kekulé, the discoverer of the structure of the Benzene molecule, who was dozing by the fire and dreaming of self-biting snakes when he made his major discovery (Boden 1990)). One of these psychology theories (Guilford 1967) concentrates on the idea of *divergent thinking*. Arthur Koestler (Koestler 1964) also wrote about a related phenomenon, naming it *bisociation*. From the computer science point of view, the modelling of divergent thinking and bisociation seems extremely difficult mainly because it is far from formally specified and, even if it was, it would certainly demand cognitive capacities that are still not achievable by computers. Yet, this does not mean that it is impossible to build models, perhaps less ambitious ones, that are capable of achieving a smaller degree of divergence, in which a computer is able to reason in a multi-domain knowledge base, eventually solving problems via transferring knowledge from different domains. Since different domains will contain different knowledge and possibly different symbols and representations, a model for reasoning in a multi-domain environment must have *translation* mechanisms, so that the transferred knowledge will still have meaning in the new context. There are well known cognitive mechanisms that establish cross-domain relationships, namely Metaphor and Analogy, which have been studied to some depth within AI, and which are certainly good candidates for plausible cross-domain transfer.

A perfect cross-domain transfer mechanism will be futile if the new knowledge is not integrated into its novel context in a meaningful way. This integration demands processes and principles able to generate knowledge structures that can be considered as a whole rather than the sum of its parts. In other words, the transfer of new knowledge should not be condemned to result in a pastiche or a concatenation of the parts, instead an emergence of new structure, function or behavior is to be favoured. Two research trends from Cognitive Science aim to solve this problem, namely Conceptual Combination and Conceptual Blending (also known as Conceptual Integration). The former traditionally deals with the understanding of linguistic combinations (such as "pet fish" or "mountain stream") while the latter is conceived as a process that can apply across the cognitive spectrum in general. Despite their differences, they both share the intent of understanding the cognitive ability of integrating knowledge from distinct sources. Both have already been subject to computational modelling.

Finally, the unavoidable question of evaluation could justify a research programme on its own, with worries regarding expertise, intentionality, complexity, aesthetic judgement, constraint satisfaction and novelty, to name only a few topics. In the current context, the evaluation should be primarily concerned with whether the just created knowledge structures are worth considering for further use and treatment within the domain it was designed for. In other words, if it is both *novel* and *useful* within this domain. The computational approach to novelty assessment has been based on similarity metrics or clustering techniques while determining usefulness is normally done via application of rules or checking constraint satisfaction. Conceptual Blending proposes a set of generic *Optimality Constraints* that aim to govern the generation of a blend. However, these are not explained formally, raising the challenge of their computational modelling.

We have just summarized some of the components for a *Model of Concept Invention* from cross-domain transfer. By concept invention, we mean the generation and addition of a new concept (or knowledge structure) to a domain in which this new concept could not be obtained by any internal means (e.g. deduction) and which can still be accepted as a valid concept for the domain. For example, before the invention of the airplane, the domain of "transport means" did not have the entire knowledge to lead to it. It was necessary to observe physical laws that were not taken into account for any other previous means of transport (even the balloon) in order to create different concepts of propulsion and sustaining structure.

3. Contributions

The main expected contributions of this book are:

- A reflection, overview and state-of-the-art survey about creativity research, according to different perspectives such as Philosophy, Psychology, Cognitive Science and Computer Science;
- A formally specified Model of Concept Invention, based on processes and principles that are coherent with the current research on creativity;
- An implemented system, Divago, which partially instantiates the Model of Concept Invention. Divago was applied to different domains and demon-

strated to be capable of generating results that pass the criteria of creativity assessment used.
- A Computational Model of Conceptual Blending, which will become integrated within Divago. This is the first computational approach to Conceptual Blending (Fauconnier and Turner 1998) that includes all the fundamental aspects of this framework.
- An assessment of the creativity of the results and of the system. We analyze the creativity of Divago with the frameworks suggested by Ritchie (Ritchie 2001), Wiggins (Wiggins 2001, 2003) and Colton et al. (Colton, Pease and Ritchie 2001).

4. Structure

The remainder of this book is structured as follows:
- *Chapter 2* is about research on creativity. It provides the necessary background regarding theories of creativity, computational approaches to creativity and frameworks for creativity assessment. In this chapter, the reader will also find the generic guidelines that underlie the rest of the book, namely at the level of the Model of Concept Invention (chapter 4) and of assessment of the results of Divago (chapter 6).
- *Chapter 3* starts by defining what a *concept* is in the context of our work. It also defines *concept invention* and compares it with *concept formation*, two kinds of concept building processes. Working with concepts is fundamental for this book, and specifically the framework of Conceptual Blending, which is also presented in this chapter. Conceptual Combination, a related area, is also presented, with particular focus to C^3, a system that will later on (in chapter 6) be compared to Divago. The chapter ends with an overview of computational approaches to Metaphor and Analogy, which deal with concept networks and from which we developed a part of Divago (the Mapper). After this chapter, the reader will have obtained a first insight on the practical aspects involved in this work (in chapter 5) and a clearer impression of the necessary notions regarding concepts.
- *Chapter 4* is dedicated to the description and formalization of the Model of Concept Invention. There, the reader will find a theoretical model, in the sense that it has not been totally implemented or specified up to the detail of computational implementation. This model provides a set of modules that we argue should be present in a system that is meant to invent concepts.

- *Chapter 5* describes Divago in detail, a system that partially implements the model presented in chapter 4. This description will take into account the modules of that model (with redefinition of the formalizations when necessary) and the framework of Conceptual Blending, which is the basis for the *bisociation mechanism* of Divago. After finishing this chapter, the reader will know Divago in depth, namely its knowledge representation, search mechanism and blending model.
- *Chapter 6* is dedicated to the experiments made with Divago. We show the five different experiments made: house-boat, for analysis of the search space; horse-bird, for the study of behavior of Divago with regard to the *Optimality Constraints*; noun-noun, for the generation of noun noun combinations and comparison to C^3; creatures generation, for the testing of Divago as an engine for generating concepts in a game environment; and established blending examples, for the validation of the Blending model implemented in Divago. The reader will get an idea of the behavior of the system within these different situations, namely with attention to the *novelty* and *usefulness* of the results. Throughout this chapter, we will analyze the system according to the frameworks of Ritchie (Ritchie 2001) and Colton et al. (Colton, Pease and Ritchie 2001).
- *Chapter 7* is dedicated to the final conclusions and discussion of future directions to take. There, the interested reader will find a multitude of research directions, some related to the generic aspects of creativity, concept invention and Blending, some more specifically directed towards the future developments of Divago.

Chapter 2
Creativity

In the first half of this chapter, we present approaches to creativity within Psychology, Philosophy, Cognitive Science and AI. The second half is specifically dedicated to the area of computational creativity, where we will show the state-of-the-art both at the level of the theoretical foundations and at the level of implementations.

1. Creativity Theories

Creativity has been the motivation for many lines of writing throughout human history, for it is such an appealing and mysterious aspect of our existence. However, it is also noticeable that its study, from a scientific perspective, has been neglected until the second half of the twentieth century (Albert and Runco 1999). The early twentieth century scientific schools of psychology, such as structuralism, functionalism and behaviorism, devoted practically no resources at all to the study of creativity (Sternberg and Lubart 1999). The oft-cited foremost turning point was when Joy Paul Guilford, in his *American Psychological Association* presidential address, challenged psychologists to pay attention to what he found to be a neglected but extremely important attribute, namely, creativity (Guilford 1950). The so-called first *golden age* of creativity then took place, with many newly founded research institutions. However this revolution did not last for long. In fact, from 1975 to 1994, only about 0.5% of the articles indexed in *Psychological Abstracts* concerned creativity (Sternberg and Lubart 1999). Today, it seems that the subject has gained another burst of output (the second *golden age*). Indeed, unprecedented resources are being directed towards creativity research in many areas.

In the following sections, the reader will be introduced to some of the works on creativity that influenced this book. These works come from the areas of Psychology (section 1.1), Philosophy (section 1.2) and Cognitive Science (1.4). Without having a direct influence on our own work, the contribution of Csikszentmihalyi (Csikszentmihalyi 1996) is also presented in section 1.3 for three special reasons: it is often cited in works of Creativity and AI[2]; it is also a respected and referential work within the area of Psychology; it reasserts some of the conclusions given by the previous sections, attesting their current acceptance. In section 1.5, we will complete the state-of-the-art of creativity with an overview of other works. Finally, a synthesis will be made in section 1.6, with particular emphasis on the aspects relevant to this book.

1.1. Divergent Production

Until J. P. Guilford introduced the operation of *divergent production* in his *Structure of Intellect* (SOI), creativity was generally considered a phenomenon separated from intelligence, a *state of mind* that was common for those considered gifted and a blessing for those we perceive of as lucky. Until Guilford, the prominent works could be roughly summarized to three: Catherine Cox (Cox 1926), who argued that creativity was a complex, multivariate behavior (as opposed to a single ability or trait)[3]; Helmholtz (Von Helmholtz 1896) and Wallas (Wallas 1926), the latter two being the creators of the four steps model[4] (preparation, incubation, illumination and verification) that became the classical stance, within Psychology, of what a creativity model should involve. This model is not contradicted by Guilford and still concurs with many current views of the subject. In section 1.3, we will take a closer look to this model. For the moment, we are interested in giving the reader a short overview of SOI, with particular attention to divergent production, the operation most linked with creative production.

The major aim of SOI was to give to "the concept of 'intelligence' a firm, comprehensive, and systematic theoretical foundation" (Guilford 1967). This very ambitious goal must be viewed from a historical perspective: during the first part of the twentieth century, many measuring tests of mental ability appeared, often motivated by the need to quantify "intelligence". This need was increased by the advents of the first and second world wars, when fast and effective processes of selection were fundamental for recruitment (mainly in areas such as the air force or intelligence services). The overstated relevance of testing the concept of intelligence justified the sarcastic sentence of E. G. Boring: "... intelligence as a measurable capacity must at the start be defined as the capacity to do well in an intelligence test" (Boring 1923). Guilford himself, who also made significant contributions to this area of psychometry, pointed out the lack of a coherent psychological theory behind tests in