

Ana-Maria Oltețeanu

Cognition and the Creative Machine

Cognitive AI for Creative Problem
Solving



Springer

Cognition and the Creative Machine

Ana-Maria Oltețeanu

Cognition and the Creative Machine

Cognitive AI for Creative Problem Solving



Springer

Ana-Maria Olteţeanu
Department of Mathematics and Computer Science
Freie Universität Berlin
Berlin, Germany

ISBN 978-3-030-30321-1 ISBN 978-3-030-30322-8 (eBook)
<https://doi.org/10.1007/978-3-030-30322-8>

© Springer Nature Switzerland AG 2020

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

“Insight is suddenly seeing the problem in a new way, connecting the problem to another relevant problem/solution pair, releasing past experiences that are blocking the solution, or seeing the problem in a larger, coherent context.”

– Sternberg and Davidson

*In memory of Aglaia Leonte, my extraordinary Nana,
and for Ștefan Smărândoiu, my maths teacher. Your set of values
combined with your belief in me made seeing myself through your
eyes both the highest motivation and the biggest self-indulgence.
My work is dedicated to you, as I am part of your work.*

Preface

You would think that a scientist researching creativity would have an easier task than others to trace the origins of her ideas, through professionally trained awareness of the topic. Despite this, I am not quite sure how the idea to direct my second doctoral thesis towards creative problem solving came about. I believe the topic sneaked up on me, and when I realized it, it had already blossomed.

With my first doctorate thesis being on aesthetics and classical music, I often had colleagues and friends asking me whether I was planning to produce some mix of my previous and current field of interest, and make for example music performing AI agents. But, in reality, I saw my interests in music on one hand, and AI & cognitive systems on the other, as having their root in a different place; in a curiosity about how minds work, and how they make sense of the world. My first thesis was, in essence, about how a listener could make different types of sense of the same piece of music, how the interpreter could direct this sense making as an aesthetic and emotional experience for the listener, and how the form of the piece of music supported these processes.

The second thesis became about how one can give new and creative meaning to objects, when needing to solve a problem one has never encountered before, or did not have the resources to solve. And how one can change the meaning and form of the entire problem, when that is the only way to solve it. About what kind of mind it takes to do that, and how it is or can be done.

In retrospect, perhaps it is all about how form supports process: how the form of the music piece, which one can see in the music sheet, supports cognitive processes of meaning making; how knowledge organization in a natural or artificial mind could be designed to support varied types of creative processes.

The effects of conducting the research on these topics, which continues in my lab today, were multiple and long lasting. One of them is a renewed understanding of just how great some of the cognitive capacities we take for granted are.

We tend to believe that problems have predefined solutions, that we must or can find. I now believe this is an illusion. Many problems do not have solutions. We define these solutions, and we also define the initial thing as being a problem. It is all part of the way we organize and sort through our world. And sometimes, about how we sort things in your cognitive worlds. The reason we sometimes lose respect for this is that we don't catch our creative problem solving process in action, and we often rely on solutions (and problem definitions) we have already memorized in order to act and understand things. But most of these solutions and problem definitions were invented at some point. And we can invent other solutions, and other definitions. We are problem definers and creative problem solvers, and these capacities are a marvellous thing. This is how we change our world, both externally, through innovation and action, and internally, through seeing things in different ways and integrating new meanings.

During the process of exploring these topics I had a lot of fun. I also learned a lot about tolerating ambiguity; and about creatively building tools myself, when I needed them for this research and they weren't anywhere to be found. If there is something I wish to you, reader, is that you share a little bit of the fun; also that you share in that sense of amazement at how great our capacity of making sense of the world and creating solutions is. I hope you go back to your life looking at this beautiful instrument you possess with a renewed sense of play. And always remember, especially in dark moments, that this capacity is always within you. And when you are stuck, you may be closer to insight than you think. You are a creator of worlds.

Now start reading so that I can thank my colleagues and collaborators.

When working on a thesis, one is supposed to accumulate knowledge, skills and abilities. If lucky, one also accumulates a debt of gratitude to the many people who have served to shape, challenge and encourage one's ideas to turn into science. It is my belief that this debt of gratitude cannot be expressed in words, and can only be paid forward, by honouring what one has learned in these years through what one becomes. However, a warm thank you goes to the following:

Christian Freksa for being a wonderful thesis advisor and head of group, from whose vision and character one cannot ever learn enough. We have had many insightful and interesting conversations, and I am sure many more are still to come.

Zoe Falomir for being a great colleague, artificial intelligence scientist and female role-model. For being enthusiastic about new ideas and nudging me gently to bring them to life.

Holger Schultheis, who kept an open door and, due to spatial proximity at the time, served as bouncing board for many creative ideas. For providing good criticism and advice on general empirical psychology practice.

Thomas Barkowsky, who served as a thorough and supportive reviewer at my graduate seminars, where some of these ideas were first explored. Also, his

invaluable work as a scientific manager of the SFB kept everything running smoothly, so that we could all do our work in peace.

Aaron Sloman, for interesting discussions and his inspiring ideas on analogical representation. The proofreaders of this work and the peer-reviewers of scientific articles I published which cover some of this content. Your comments have all helped improve this.

My partner Anton Mykell Sykes for making endless amounts of coffee, for understanding so many moments of random quietness or unexplainable bouts of excitement when I was caught by a new set of ideas, and for getting more knowledgeable on cognitive science and artificial intelligence every day by listening to me talk about my interests, work and process.

Finally, to the German Research Foundation (DFG) for the SFB/TR 8 that I was so lucky to be part of. This provided generous funding and support for summer schools, conferences, and a lovely place to work, think and disseminate ideas in.

Berlin,
October 2018

*Dr. Dr. Ana-Maria
Oltețeanu*

Contents

1	Introduction	1
1.1	Book Structure	4
<hr/>		
Part I Rebraiding the Strands: Creativity and Problem Solving, Human and Computational		
<hr/>		
2	Creativity, Problem Solving and Insight	11
2.1	Creativity	11
2.2	Problem Solving	14
2.3	Insight	16
2.4	Neural Correlates of Insight	19
2.5	Visuospatial Intelligence	20
3	Knowledge Organization for Creative Problem Solving	25
3.1	Representations and Structure	27
3.2	Creative Processes	29
3.3	Restructuring or Re-Representation and the Link between Knowledge Representation and Process	31
4	Computational Creativity and Systems	33
4.1	AM	34
4.2	HR	34
4.3	Aaron	34
4.4	The Painting fool	35
4.5	Poetry systems	36
4.6	BACON	37
4.7	Magic trick making	37

5	Two Types of Evaluation: Human Creativity, Computational Creativity	41
5.1	Creativity Testing for Human Participants	41
5.2	Creativity Evaluation in Computational Creativity Systems	48
5.3	Summary of Part I	53

Part II One Ring to Creatively Solve Them All? In Search of a Unified Cognitive Framework

6	Gathering the Cogs of a Framework	59
6.1	Creative Problem Solving Redefined	60
6.2	Desiderata for a Framework of Creative Problem Solving	62
6.2.1	Visuospatial Inference	62
6.2.2	Creative Use of Affordance	63
6.2.3	Concept Generation and Structure Transfer	64
6.2.4	Insight and Re-Representation	67
7	CreaCogs – A Clockwork View of Creativity	69
7.1	A Clockwork Metaphor	69
7.2	The Clockwork	81
7.3	Desiderata through CreaCogs’s Lens	84
7.3.1	Visuospatial Inference in CreaCogs	84
7.3.2	Creative Use of Affordance in CreaCogs	86
7.3.3	Concept Generation and Structure Transfer in CreaCogs	86
7.3.4	Re-Representation through Search and Construction in CreaCogs	88
8	The Cogs in Motion – Navigating the Knowledge in CreaCogs	91
8.1	Feature Spaces	91
8.2	Concepts	92
8.3	Problem Templates	94
8.4	General Creative Mechanisms	96
8.4.1	Hypothesizing by Concept Similarity	96
8.4.2	Creative Use of and Hypothesizing with Problem Templates	97
8.5	Creative Search and Substitution Mechanisms	98
8.5.1	Upward Search	99
8.5.2	Downward Search	101
8.5.3	Sideways Search	101
8.5.4	Restructuring, Creation of New Objects and Re-Representation	102
8.5.5	Discussion	103
8.6	Creative Construction Mechanisms	104

8.7	Discussion – the Relation between Imperfect Recognition, Constructive Memory and Re-Representation	106
-----	---	-----


Part III Empirical and Computational Explorations

9	What Do Swiss, Cake and Cottage Have in Common? – Computational Explorations of the Remote Associates Test	113
9.1	Meet the RAT (The Remote Associates Test)	115
9.2	Which Cogs Fit the RAT?	116
9.3	A Computational RAT Solver (comRAT)	118
9.3.1	The RAT Knowledge Base	118
9.3.2	Knowledge Acquisition and Organization by Association	119
9.3.3	Solving the Compound RAT (cRAT) Queries	120
9.4	Formalizing Preferred Answers and Item Contribution	123
9.5	comRAT-C Experimentation and Results	125
9.6	Preferred Convergence – Empirical Analysis	127
9.7	RAT Difficulty for comRAT and Humans	129
9.8	What Does It Take to Be a RAT? – Generative Abilities of comRAT	131
9.9	Not a Blind RAT – Expanding the Remote Associates Test to the Visual Domain	133
9.9.1	Making the RAT Visual	134
9.9.2	Giving the Visual RAT to Human Participants	135
9.9.3	Solving the Visual RAT Computationally – from comRAT to comRAT-V	137
9.9.4	Human Participants Results When Solving the Visual RAT	138
9.9.5	Results with comRAT-V	140
9.10	Discussion and Further Work	142
9.11	Conclusions	147
10	What Could You Use This Object for? – Object Replacement and Object Composition in a Computational System, and the Alternative Uses Test	149
10.1	Cognitive Knowledge Retrieval and Encoding	150
10.1.1	Multi-Feature Object Recognition	150
10.1.2	Multi-Feature Similarity	153
10.1.3	The Various Features Spaces Obtained	155
10.2	OROC Implementation	156
10.2.1	Object Replacement Question Types and Proof of Concept	157
10.2.2	Object Composition Question Types and Proof of Concept	160
10.3	Evaluation of Performance with the Alternative Uses Test	162


- 10.3.1 Constructing an Evaluation Procedure 162
 - 10.3.2 Experimentation and Evaluation with Human Judges .. 164
 - 10.3.3 Evaluation Using Human Think Aloud Protocols 168
 - 10.3.4 Human Alternative Uses Answers 169
- 10.4 Discussion 171
- 11 Daily Eureka's About Candles and Strings. Approaching Insight with Practical Insight Problems 175**
 - 11.1 Setting up the Practical Insight Problems 177
 - 11.1.1 Classic Practical Object Insight Problems 178
 - 11.1.2 Strategies for Creating New Problems 186
 - 11.2 Encoding Strategy and Codes 192
 - 11.3 Two Cases Studies 203
 - 11.3.1 The Case of Mb8 204
 - 11.3.2 The Case of Fb6 211
 - 11.4 Analysis of Process and Codes in the Context of CreaCogs.... 218
 - 11.5 Order of Process for Modeling in CreaCogs 233
 - 11.6 Towards Modeling Practical Object Insight Problems in CreaCogs 237
 - 11.6.1 Modeling the Candle Problem..... 237
 - 11.6.2 Modeling the Two Strings Problem 239
 - 11.7 Discussion and Future Work 241
- 12 The Journey Thus Far and the Journey Ahead 245**
 - 12.1 The Journey Thus Far – Contributions Overview 245
 - 12.2 Experimental Work and the CreaCogs Framework 246
 - 12.3 A Comparability Approach to Creative Computational Cognition 249
 - 12.4 The Journey Ahead 250
- A Appendix A – Human and System Answers to Variants of the Remote Associates Test 253**
- B Appendix B – OROC Results 265**
- References 271**

Introduction

The story goes that Kekulé day-dreamt of an Ouroboros symbol (a snake eating its own tail) or a Tibetan knot, when trying to find the structure of the benzene molecule (Fig. 1.1a). An equally mesmerizing tale tells of Watson dreaming of spiral staircases before doing his part in coming up with the double helix structure of DNA (Fig. 1.1b).



(a)



(b)

Fig. 1.1: Dreams and day-dreams: (a) Kekulé's, (b) Watson's.

One can easily see the visual similarity between the (day)dreamt object and the discovered “object”. Unfortunately we cannot take such stories at face value, no matter how beautiful. Nor can we engineer for people to discover new molecules every day and in controlled conditions in the experimental lab, in order to study how they do it. The questions such stories trigger, though, are very alluring ones:

- how do humans solve such problems creatively? And much simpler problems for that matter too, as creativity is a thing often encountered in daily life, and
- will we ever be able to create artificial cognitive systems that do the same, or that have such insight into cognitive processes, that they can help humans creatively problem solve more often and with more ease?

This book is an investigation of these questions and offers the beginning of a possible answer.

The subjects of creative problem solving and productive cognition addressed in tandem are topics of great interest to both Artificial Intelligence and Cognitive Science. For Cognitive Science, designing systems that can perform various levels of creative problem solving and testing hypotheses on such systems can contribute to the proposal of further cognitive models of creativity, thus helping us understand how the human mind works when performing various types of creative problem solving. For Artificial Intelligence, the topic can set the foundations to enable the next generation of creative assistive systems — systems which can make creative associations and propose novel solutions or new lines of enquiry; these creative inferences should be expressed in ways which are easy to comprehend by humans and integrate in a normal workflow; the more such systems understand how human creativity works, the more they could provide cognitive support for it. AI and CogSci working together have historically yielded many great achievements and ways of looking at the core questions of what a mind is and what does it take to make one.

For Computer Science, the topic can set to explain how new (valid) information can be created out of old information, other than by pure logical inference. Finally, for Philosophy of Information (Floridi, 2011), the topic might help us define the limits of generative systems, and new ways of measuring informativity.

It is thus worth endeavouring to rebraid together the topics of creativity and problem solving on one hand, and those of computational and human cognitive processing on the other. Creativity (Boden, 2003) has been studied lately in more computational terms, with many computational creativity (Colton & Wiggins, 2012) systems being created. Initially, most such systems aimed to implement artistic endeavours — like poetry writing (Colton, Goodwin, & Veale, 2012) and painting (Colton, 2012b), with only a few tackling creative problem solving (Fleuriot, Maclean, Smaill, & Winterstein, 2014; Bou et al., 2015). Though a move against systems which merely generate artifacts has been made (Ventura, 2016), and the field is deeply preoccupied with the issue of evaluation, most such endeavours aim to enable computational cre-

ativity systems, to ask to ask what it is to be computationally creative, and how can it be evaluated.

That is to say, most of these systems do not set out to account for the cognitive mechanisms producing these creative results, and are merely inspired by the results of the creative mind rather than trying to elucidate its processes. Thus, the differences between various creative processes are not accounted for, the special status of some such processes — like insight (Batchelder & Alexander, 2012) — are not investigated in the computational creativity community (with the exception of concept blending (Fauconnier & Turner, 1998) and metaphor (Veale & Keane, 1992)), but mostly in the cognitive psychology literature. Psychological hypotheses on the stages of such processes are not taken into consideration, thus no further elaboration and investigation of these stages comes as a result of designing such systems. Nor are such systems able to be tested with the same tests which we give to humans (Duncker, 1945; Maier, 1931; Mednick & Mednick, 1971). On the other hand, theories of creative cognition are not implemented as often as they could be (with some exceptions like analogy (Falkenhainer, Forbus, & Gentner, 1989; Hofstadter, Mitchell, et al., 1994) and incubation (Hélie & Sun, 2010)), nor are they implemented in a unified manner — with one implementation or small set of principles acting as one architecture through which multiple creativity processes can be modeled and explained. However, new tools from computational creativity might allow the implementation of many more such theories, if only such tools would endeavour to take cognitive processing, cognitive knowledge acquisition (the kind of knowledge humans have) and knowledge organization into account.

To fill this gap between cognitive psychology, AI methods and the new field of computational creativity, this work aims to design hypotheses and implement systems which are:

- a) in line with existing work in the cognitive science literature;
- b) at levels of description which are adequate for cognitive science (discussing possible representation and processes) and
- c) on which further cognitive models can be developed, and empirical hypotheses of how such creative processes work can be explored and tested.

Furthermore, some of these systems and the hypotheses on which they are constructed are evaluated using tests given to humans and products of human creativity as a comparison. Some others, as a result of this work, are put in a form which allows such evaluation in future work.

A main hypothesis of the following work is that knowledge organization is a key factor when approaching creative problem solving. We posit thus that knowledge organization approaches which can naturally and with ease support creative processes in computational systems need to be designed and refined. Throughout this work, knowledge organization is approached and implemented in ways which enable creative search, re-representation of previous knowledge, combinatorial creativity, associativity with similar terms and convergence upon solutions.

This is an issue of knowledge organization, not knowledge representation. The forms of representation chosen in the framework and implemented in the various systems here can be changed while maintaining similar results. However, the organization of said representations is a core principle, enabling the creative process to happen without high computational costs, in the same way in which various data structures are better at dealing with and representing various processes.

1.1 Book Structure

We will start this journey into discovering what it takes to be or build a creative problem solving cognitive system by rebraiding the strands of creativity and problem solving, in the realms of human and computational skills. We will thus look at various research threads: first at creativity and problem solving from the perspective of human creative cognition; then at problem solving from the perspective of artificial intelligence.

It is the thesis of this book that particular types of knowledge organization will have higher chances of enabling (the implementation of) creative processes. Because of this, the types of processes which have been proposed to enable creativity or problem solving and the types of knowledge which could support such processes will be briefly reviewed.

Various computational creativity systems have been implemented in the last decades. A short tour of a selection of such systems will give a taster of what creative machines the community is building in domains as varied as mathematics and magic trick making. We will then return to integrating the human and the computational, by exploring how the forms of evaluation of computational creativity systems compare to those used to assess creativity when dealing with human cognition.

It is all well and good to aim towards an integrated view of how natural and artificial cognitive systems can problem solve creatively. However, how diverse would the requirements for such a creative problem solving system be? And do we have any chance of ever addressing such a diversity of tasks with a small set of processes? In Part II we have a look at a subset of diverse requirements and put together a framework; this theoretical framework proposes a type of knowledge organization and a small set of processes aimed at solving a diverse number of creativity tasks. These processes and the framework are then explored and partially formalized.

To put some of these principles to the test and have some hands-on fun, in Part III we proceed to empirically and computationally explore them, in experiments involving both human participants and programs, with the help of our previously defined framework. First, we explore what Swiss, Cake and Cottage have in common — that is how humans solve the Remote Associates creativity test (RAT). A computational system which can solve the same task is implemented and experimented with, to test the associative principles of our

framework, explore one of the earlier proposed mechanisms and to understand more about how humans solve such problems. Helped by this system, and the new understanding of the task, we ask what does it take to be a RAT — thus what are the principles the Remote Associates test is based on. We then construct a visual version of this creativity test, based on the same principles of Remote Associates, and give it to humans to solve.

Us, humans, generally ask computational systems to solve problems we have created or stumbled upon. To reverse the roles for a bit, we proceed on making our computational RAT solver create queries, rather than solve them. More than a fun pursuit, this will help us give cognitive psychologists sets of RAT queries in which they can control many more of the variables, thus allowing them to study the processes involved in solving such queries with a higher degree of precision.

We further ask what you could use a cup for. This seemingly inconspicuous question hides a wonderful skill: the ability of using objects creatively, which most humans and some animals seem to have. We build a prototype system that recommends to use dental floss if you are missing a clothesline. The system (OROC) tests some other principles of the theoretical framework proposed above, by doing creative object replacement (OR), and also a bit of object composition (OC). We pair this prototype with the Alternative Uses test, used to test creativity in humans. Will OROC be able to give sensible answers? We let humans judge its skill, evaluating it with the same metrics as humans solving the Alternative Uses test would be evaluated with. Amongst others, we investigate what kind of answers such a system gives, compared to humans, if its process is in any way similar to that of humans and what kind of properties make humans think the answer is a particularly good one.

Towards the end we finally touch upon the higher level issue of insight. Insight capable computational solvers would require a large amount of knowledge to be built in, to even begin to test quirky creative processes. We prepare for this by using the previous experiments as a gateway, and approaching insight problems in a domain in which we will be soon prepared to implement solvers.

Have you ever wondered how an insight problem is created? Insight moments might seem like aloof and legendary moments of discovery; however, empirical insight is studied in the lab, and insight problems for these settings need to be created (by not so aloof or legendary humans, though see (Duncker, 1945), at least on the legendary). We put together a strategy for creating such insight problems, based on our reverse-engineering of some of the processes we understand to be implicated in insight. We then give these and classical insight problems to people to solve, in a think aloud protocol — with people speaking as they solve each problem. We come up with a set of codes to classify and compare their answers to our theoretical framework. This is to help us explore whether the same framework principles posited before could indeed or could not be applied at the insight level.

An intriguing and at times quirky journey awaits us, which needs an open mind. Creative problem solving answers might not be perfect answers, but they reinvent our way of seeing the world, or our way of understanding a set of matters which was previously ambiguous. Before reaching insight, though, let us start at the beginning. A lot of inspired research has gone into creativity and problem solving, from different angles of interest: sometimes on both creativity and problem solving, sometimes on each separately; at times from the computational perspective, and at others from the human cognition perspective; finally, sometimes (though a bit rarer than we would like), such investigations have happened in the interdisciplinary spirit of cognitive science. Though we cannot hope to honour all the work that happened before us, in the following pages we will get a taste of it, and hope to leave it slightly better off at the end.

**Rebraiding the Strands: Creativity and
Problem Solving, Human and Computational**

Problem solving and creativity are often addressed together (Sternberg & Grigorenko, 2003) as higher level cognitive abilities. Both have been held in high esteem and long considered to be human-only abilities, and then proven to exist to a smaller yet still impressive extent in animals: other animals are capable of some creative tool use (Köhler, 1976) and analogy-making (Gillan, Premack, & Woodruff, 1981), and frameworks for the study of animal creativity have been proposed (Kaufman & Kaufman, 2004; Bailey, McDaniel, & Thomas, 2007). However, creativity and creative problem solving are at their pinnacle in human cognition.

Extraordinary leaps of thought have been an integral part of human history: one only needs to leaf through Haven's or Philbin's lists of greatest inventions of all time (Haven, 2006; Philbin, 2005), or through Watson's history of thought and invention (Watson, 2005, 2011) to reinstate in one's self a sense of awe regarding the human ability for creative thought. From the windmill to inventing impressionism, such inventions, brought forth by individuals or groups, seem to be leaps of thought. Nonetheless, creativity is encountered in the everyday life of most people: it happens in the kitchen, when you are producing a creative variation of a recipe, when you are repairing items around the house with improbable tools and when you are adaptively problem solving an unexpected event, creating new plans for the day.

Despite the universality and the diversity of levels creativity takes, various difficulties relating to knowledge representation, common sense knowledge and the amount of cognitive functions involved in higher level cognitive abilities stand in the way of directly modeling such processes.

Both creativity and problem solving have been addressed in psychology, AI and cognitive science, and can be conceived of as interdisciplinary fields of research. Different kinds of matters pertaining to these subjects have been studied, depending on the field doing the inquiry. Here is an example of how the field of inquiry affects the question. A question formulated by *cognitive psychology* would be: "*How does a certain creative process function in humans? How can we model it?*". An *Artificial Intelligence* type of question: "*How can we define problem solving so that it is computable by machines and that computation can be optimized?*". Questions relating to *cognitive science*: "*What kinds of representations and processes are necessary and sufficient to have creativity in a cognitive system?*". An important question asked by the newly emerging field of *computational creativity* is "*What are the required criteria to call a system creative, and how can one evaluate such creativity?*". This question seems to come loaded with the a bias shown by recent research: as soon as humans *know* by which process the system is generating its creative products, they become much more reluctant to call it creative, as if removing some of the mystery about the process might make it less valuable to some people. As the history of thought regarding human creativity comes pre-laden with such mysterious concepts as muses, daemons and inspiration, it is no wonder that peeking into the process might make some people feel that an essential aspect of creativity has vanished.

However, for the rest of us, for whom the state of wonder remains unaffected or grows in the advent of peeking behind the curtain of cognitive processes, a systematic synthesis is needed to understand the interrelations of the various fields, and the work that has gone before. To this purpose, a large amount of literature from these various domains is approached systematically in the following chapters, by grouping these subjects along these lines:

- **Chapter 2 – Creativity, Problem Solving and Insight** – describes relevant theories of creativity, problem solving (in its well-structured and ill-structured forms) and insight, as an empirically studied creative problem-solving process. Do sight and insight bear a relation? The chapter concludes by exploring which aspects of visuospatial intelligence one needs to keep an eye on or understand the structure of when talking about creative problem solving.
- **Chapter 3 – Knowledge Organization for Creative Problem Solving** – addresses the question of what kinds of representation and processes are of specific use when attempting to implement and model creative problem solving. Various representations and processes previously considered relevant to creative problem solving are reviewed. The interplay between representation and process is presented as a motivator for searching for and engineering types of knowledge organization which can support the creative process in its most relevant forms.
- **Chapter 4 – Computational Creativity Systems** – reviews formal and applied work on creative systems, presenting a selection of models and computational creativity achievements. From computational painters to computational magicians, this review can only offer a selective taster of the field, and many other interesting systems have been realized in the last years.
- **Chapter 5 – Evaluation of Human and Computational Creativity** In this section, modern work and thought on the topic of evaluation for computational creativity systems is reviewed side by side with psychology work on the assessment of creativity in human participants. This juxtaposition is meant to enable a fruitful comparison between the ways in which evaluating humans and evaluating artificial systems has been done so far.



Creativity, Problem Solving and Insight

What are creativity, problem solving, and insight? This section sets to present the conceptual work various researchers have put into defining these terms.

2.1 Creativity

Various theories of creativity exist, addressing various aspects of creativity. For example, a distinction is drawn between inventing the bicycle, and coming up with the creative thought that you could use a shoe to put a nail in the wall. The first is called **historical creativity** (h-creativity) by (Boden, 2003), representing creative acts which are original on the scale of human history. The second, though personally quite satisfying, is called **psychological creativity** (p-creativity) in her taxonomy, and refers to contributions which are creative from the perspective of the individual.

Boden further differentiates between **combinatorial**, **exploratory** and **transformational** creativity. Combinatorial creativity is a form of producing new, unusual combinations or associations out of known ideas. Boden's example is that of a physicist comparing an atom to the solar system. Exploratory creativity is a process of exploration of variations within a certain conceptual space. In Boden's words, "*Within a given conceptual space, many thoughts are possible, only some of which may actually have been thought. [...] someone who comes up with a new idea within that thinking style is being creative in the second, exploratory sense.*" Transformational creativity is about changes to/restructuring of the conceptual space altogether. Boden exemplifies it as "*someone thinking something which, with respect to the conceptual spaces in their minds, they couldn't have thought before. [...] (the preexisting style) must be tweaked, or even radically transformed, so that thoughts are now possible which previously (within the untransformed space) were literally inconceivable*". The term of *conceptual space* here is, according to some (Ritchie, 2001; Wiggins, 2001), vaguely defined. It is then hard to compare it to similar

terms in the literature, for example the conceptual spaces used by Gärdenfors (Gärdenfors, 2004).

Margaret Boden is the author of some of the most well cited recent thought on creativity, however thought on creativity is by no means new. Take another term relevant for the study of creativity, that of **divergent thought**. Currently, divergent thinking implies a method of generating creative ideas by exploring many possible solutions. Guilford came up with the term *divergent production* in the Structure of Intellect (Guilford, 1967), referring to a form of broad search for many and varied solutions; such a search would happen mostly spontaneously and in a free flowing manner. Convergent thinking, coined as its opposite, complementary cognitive capacity, stood in his theory for a process of coming up with a solution by following a set of logical steps. The solution generated convergently would bear many restrictions and be rigorously structured. Convergent and divergent productions, the differences of which can be observed in Table 2.1, are considered by him as complementary parts of the cognitive *productive* capacity. Guilford’s model for convergent and divergent productions is a common model, arguing that problem solving and creative production are the same.

Table 2.1: Difference between divergent and convergent productions, according to Guilford

Divergent productions	Convergent productions
Loose and broad problem, or incomplete grasp of it at the agent level	Answer can be rigorously structured and forthcoming
Few restrictions	Many restrictions
Broad search	Narrow search
Vague and lax criteria for success (stress variety and quantity)	Sharper, rigorous, demanding criteria

The **generative-exploratory** model or Geneplore by (Finke, Ward, & Smith, 1992) differentiates between two phases of creative thought: generation and exploration. In the generative phase, preinventive structures (which are mental representations) are constructed by the individual. In the exploratory phase, these structures are used to generate new ideas. This model, like many others, has been criticised as being too vague to implement in a program.

Gabora has proposed creativity to be a **honing** process (Gabora, 2005; Aerts & Gabora, 2005). In her view, the concept of an individual world-view takes center stage: creativity is then a process by which an individual hones, at multiple stages, their world view. This honing of a worldview is a

self-organization process which aims to solve inconsistencies between ideas, attitudes and knowledge.

Other models of creativity center on the difference between implicit and explicit knowledge. **Explicit** knowledge is knowledge which can be articulated, communicated, recorded and distributed (words, numbers, mathematical and scientific formulae). **Implicit** knowledge is the opposite: knowledge which is not easy to communicate, or knowledge which the knowledge holder might even be unaware they possess. Trying to explain to someone how you ride a bicycle might give you flavour of a type of implicit knowledge: such knowledge is hard to express because it is sensorimotor in origin, rather than verbal (for some, such knowledge is called procedural).

Implicit knowledge can also refer to knowledge that you do not know you have acquired, nor can you consciously recall or recollect. If you cannot recollect it, how is such knowledge shown to exist? The presence of such knowledge has been shown in experiments on priming and skill learning (Schacter, Chiu, & Ochsner, 1993). Examples of priming involve being given word fragments with multiple possible completions (e.g. a-a-in to be turned into assassin) and completing them with previously studied items, which are however not remembered; or being given sets of lexical terms like “flig” and asked to decide whether they are words or non-words, and making the decision faster for previously studied items, though these items are not remembered. Such priming experiments use linguistic stimuli, auditory word stimuli and visual stimuli (pictures). In terms of skill learning, classic studies (Milner, Corkin, & Teuber, 1968; Glisky & Schacter, 1988) have shown that amnesic patients can acquire new perceptual and motor skills.

Creativity has historically harboured an aura of mystery, and sometimes deals with ideas and solutions which appear fully formed through a flash of insight. It is easy to see how implicit processing – that is processing which would happen without the awareness of the participant – would be an interesting topic from the studying creativity perspective. Some scholarship on creativity thus focuses on or at least integrates implicit processes. For example the **Explicit-Implicit interaction** model (EII) (Hélie & Sun, 2010) proposes a unified framework for understanding creativity in problem solving, based on the relationship between implicit and explicit processes. The EII theory has been implemented in the CLARION cognitive architecture and relies on a set of principles which include the coexistence and simultaneous involvement of implicit and explicit processes in most tasks.

Other models of creativity have also been proposed (Schmidhuber, 1991; Thaler, 2013); such models can be big picture views, and may or may not be linked to specific creative processes, like analogy, metaphor and concept blending. We will explore such types of specific processes in Chap. 3.2. However, first it is important to realize that creativity and problem solving are not often discussed in conjunction.

2.2 Problem Solving

Part of the difference between creativity and creative problem solving comes from what the results of each are, and how those results are evaluated. Thus, when speaking of creativity without the context of problem solving, one tends to generally consider the processes of creating works of art – music, poetry, paintings – which are original (the field of computational creativity is not devoid of this bias). Such works can then be evaluated in terms of their aesthetic qualities, their novelty compared to other works in the similar genre, compared to other works of the author, or their novelty in terms of process. Creative problem solving on the other hand has to satisfy problem constraints: whether the new solution has a chance at satisfying the problem matters. The usefulness of the solution can thus be as much a factor in evaluation as novelty. The field of creativity usually gets closer to the fields of problem solving and reasoning when it deals with innovation, scientific discovery and scientific reasoning (Langley, 2000; Klahr & Dunbar, 1988; Nersessian, 2008).

In order to understand creative problem solving, one must thus revisit classical problem solving definitions in Artificial Intelligence (AI), and aim to refine as to allow for the creativity component.

In AI, and later in computer science, problem solving in its classical form is defined (Newell & Simon, 1972) in certain specific terms. We will enumerate these terms, showcasing them in a classical example: the tower of Hanoi problem. As shown in Fig. 2.1, three rods and a set of disks which can slide onto the rods. The solver is supposed to move the entire stack to another rod, with the constraints that only one disk can be moved at a time, only the uppermost disk of a stack can be moved on top of another stack, and only smaller disks can be placed on other disks. Problem solving has been defined in terms of:

- An initial state of the problem. For example, the depiction in Fig. 2.1 is the initial state of the tower of Hanoi problem.
- Operators or successor functions which define reachable states of the problem $f(x)$, from any state x . A move of a disk on a larger disk or on an empty rod would constitute an operator. Thus reachable states from the initial state shown in Fig. 2.1, through one move operators, are having the red disk on either the A or the C rods.
- A state space, constituted of all the reachable states, based on applying the operators to initial states in whatever sequence. In the context of the tower of Hanoi problem, this would include all the possible states of disks in various decreasing orders on various rods that can be obtained applying the operators allowed above on the initial state of the Hanoi problem. The state space will thus include having disks 1 and 2 on rod C, as this can be achieved while respecting the available operators.
- Paths – sequences through the state space. In the tower of Hanoi, this would mean a certain set of disk moves, which will allow the navigation through different possible states of disk configurations. To achieve the

state of disks 1 and 2 on rod C, for example, the following path can be taken: 1 can be moved to A, 2 can be moved to C and then 1 can be moved on top of C.

- Path cost – a function used to evaluate the best heuristics. In tower of Hanoi, this could be the number of moves on a different move sequences (path) to reach the goal configuration.
- Goal state or goal tests (to determine if the goal state has been reached). In our example, the goal state is having all the disks on a specific other rod.
- Heuristics, which can be defined based on their success and cost (optimality). For the tower of Hanoi, a heuristic is to keep moving the smallest piece from the initial rod on a different rod, move a larger piece on a different rod, then put the smaller piece(s) on the larger piece.




Fig. 2.1: The tower of Hanoi problem

However, a difference can be drawn between well structured and ill-structured problems (Newell, 1969). Well-structured problems can easily be described in terms of the classical problem solving definition, while ill-structured problems have ambiguous initial states, operators or goal states. Take the problem of the benzene molecule – what is the goal and what are the exact operators? Ill-structured problems are more often encountered in real environments (Simon, 1974) than well-structured problems: designing a house, increasing the water supply for a growing community, making a budget or designing a cognitive systems are all ill-structured problems.

One way to address ill-structured problems might be by using **productive**, rather than **reproductive** thinking. The distinction between the two was made by the Gestalt psychologist Wertheimer (Wertheimer, 1945). He considers reproductive thinking to be a function of repetition, conditionings, habit and familiar ways of thought. This means that applying the same known routines to solve a problem could be understood as a function of reproductive thinking. For example, calculating the length of the hypotenuse in a right triangle by applying the known pythagorean theorem, or baking a recipe which you already know can be seen as an example of reproductive problem solving. Meanwhile, productive thinking is considered to produce new ideas and be

insight-based. For example, coming up with ways to calculate the length of the hypotenuse in a right triangle when not knowing the theorem, or coming up with a recipe that you can make from the ingredients at hand are examples of productive thinking. In general, coming up with new ideas and new ways of doing things which solve the problem (or, if imperfect, might solve another problem), can be seen as a function of productive thinking. Applying this definition to problem solving, productive problem solving can be defined as a process which brings about new heuristics and new ways of looking at the problem, while reproductive problem solving would mean applying the same known heuristics to the same types of problems. Productive problem solving is thus what we call creative problem solving.

In order to understand why certain ill structured problems are hard to solve and require productive problem solving, the case of insight problem solving can be taken as an example.

2.3 Insight

The legend has it that Archimedes jumped out of his bathtub shouting “*Eureka*” because of having an insight on how to measure the volume of a crown while observing himself immersed in the water (Vitruvius Pollio, 1914). Various other anecdotes about moments of insight in scientific discovery, exist, like the ones mentioned before of Watson and Kekulé. Some such anecdotes are introspective accounts, declared (some time after the actual insightful event has happened) by the solver, thus standing chances of being distorted by reporting later (due to imperfect memory of the event) or for the sake of a good narrative.

To delve deeper into the empirical study of insight (Chu & MacGregor, 2011), one cannot rely on the introspective or anecdotal account of such insight moments regarding h-creativity discoveries. For the empirical study of insight, empirical tasks do exist. Some of them are problems, like the matchstick problem shown in Fig. 2.2. In this problem you are supposed to make the equation with roman numerals true, by moving only one matchstick, without removing it from the equation. Many people attempt to solve this problem by manipulating the various quantities, for example by moving the matchstick before the first numeral (IV) after the third numeral (I), thus turning the first numeral to V and the third to II. This problem approach is assumed to happen because of functional fixedness — that is being stuck in a perspective that certain objects can be manipulated (or are functional) only in a certain way. In the problem below, such functional fixedness could derive from having learned that, in equations, it is quantities that are being manipulated. The problem is, however, solved by manipulating not the quantities, but the signs — thus removing a matchstick from the initial equal sign, and adding it to the subtraction sign. This turns the first operator into a minus sign, and the second into an equal sign, changing the equation into a correct one.