

ARTIFICIAL CREATIVITY

Pascal Wallisch
University of Chicago
Student-ID 923227
Lascap@gmx.de

- Contents -

1 Introduction.....	3
2 Artificial Intelligence: A brief overview	5
2.1. Initial works in AI and the history of AI.....	6
2.2. Problems of AI research.....	9
2.3. Outlook and applications.....	14
3 Creativity.....	17
3.1. Conceptual basics.....	17
3.2. Theories of creativity	21
3.3. Empirical issues in creativity research.....	26
3.4. Summary.....	33
4 Artificial creativity	35
4.1. Foundations and conceptions.....	35
4.2. Specific examples of computer models of creativity.....	36
4.3. Evaluation of the artificial creativity approach.....	43
5 Outlook and conclusions	48
6 References.....	50

1 Introduction

Creativity is one of the most thrilling problems of contemporary psychology. The sheer possibility of creative thought and innovation seems to lie beyond the realm of what can be explained within our framework of natural science. The common man, on the other hand seems to be comfortable with these problems and uses concepts like “intuition” and insight to explain the obvious possibility of creative achievements or attributes these performances to romantic notions of the genius that is inspired by his muses.

Of course, these conceptions of folk psychology don't withstand a thorough analysis by psychological science. Insight and intuition, used as explanatory constructs are mere labels for a process that we obviously understand very, very poorly on all levels of analysis.

For a long time, creativity research was basically done and creativity explained within an – now outdated – psychoanalytical framework. This framework didn't exceed the explanations provided by folk psychology much and was basically empirically untestable, like the other rantings of psychoanalysis, still holding interesting psychological constructs like “envy” or “hate” firmly in their grasp, dominating these fields and ruining them for everyone else.

Modern creativity research was begun anew by Guilford in the fifties, rejecting psychoanalytic theories and framing creative concepts in a way that it became a (seemingly) tractable problem.

After that, creativity research became prominent both in personality psychology, trying to explain the huge apparent differences in interindividual creativity and cognitive psychology, still trying to figure out the nature and mechanisms of creative thoughts and processes.

With time, these new approaches achieved many theoretical differentiations and empirical findings in the field of creativity research, while the overall understanding of the creative process and creativity remained surprisingly poor: A rough measure that proved to be useful for other fields of investigation (Wallisch, 2001) confirms the suspicion that we use the term far more readily in everyday life than we know about it. Psycinfo cites only 13254 articles

with the keyword creative or creativity in it, but Google lists over 8.820.000 websites that use the term, leading to an “overuse”-index of 665. This is even worse than our use of attention: For Attention, we currently find 14.500.000 websites containing the word “attention”, but there are only 67857 articles with the keyword attention on Psycinfo (index-score: 213).

For comparison, our understanding of Perception is obviously far better: 2.450.000 citations on Google, 150419 articles on Perception in Psycinfo (index-score: 16).

This usage of the concept indicates that we like to talk about it a lot, without really knowing what is going on. On top of that, a lot of the citations don't really deal with creativity itself but use the term in a more metaphoric way. For a more detailed discussion, see chapter 3.

In parallel, also starting in the fifties, the science of artificial intelligence began to address computational problems of cognition, traditionally handled by psychologists. After some initial successes, the field came to a standstill for decades, not making any significant progress, as it became apparent that the brain doesn't process information like a computer and that the computer metaphor is misleading at best. While some preliminary successes could be established, the great hopes that started this field didn't come true, which will be discussed in much more detail in chapter 2.

Creativity research is therefore still a field of great mystery and wonders. While there is much speculation in psychology about the mechanisms involved in the creative process that produce creative products, solid neurophysiological or neuropharmalogical evidence is nonexistent. These fields still consider creativity to be a non-issue. Therefore, no one can look to neurobiology for solutions of the nagging theoretical problems, solutions that could resolve the current sorry state of this field.

It was only in the last decade that the application of the principles of artificial intelligence to creativity issues brought many new and intriguing insights to the field of creativity, concerning applied as well as theoretical issues. More on this in chapter 4.

The paper concludes with chapter 5 that summarizes this short essay.

2 Artificial intelligence: A brief overview

Artificial intelligence is – even more so than creativity itself – a huge field of investigation of its own. Due to limitations in space, I will only discuss aspects of both fields that are relevant for a better understanding of artificial creativity. This entails that I have to neglect many important issues in artificial intelligence, including almost all philosophical implications of research in the field, many of which are of crucial interest for psychology in general, but only of limited relevance for our issue at hand. The content of this brief overview derives largely from Boden (1990), Winston (1992) and Thagard (1999). Please see these references for a more extensive discussion of the issues raised in this chapter.

For our purposes, a grossly simplified definition of Artificial Intelligence (AI) will be sufficient: AI can be conceptualized as a field of investigation that strives to mimic, simulate, emulate, model or even generate human information processing by man-made artificial means, mainly computers and computer programs. In this framework, the ultimate goal is therefore the construction of a “thinking” machine, thereby fulfilling another human dream.

Please note that this AI-endeavor bases on a crucial assumption, namely that humans themselves can be understood as sophisticated information processing machines and that their (problem-solving) performances can be described in computational terms. This implies, that all human performances that are considered to be “intelligent” can be – in principle – modeled by artificial computational devices. The many other implications of this critical assumption can be seen throughout the chapter.

This chapter is divided in three parts: First, we briefly discuss the history of AI, then go on by reviewing fundamental problems of AI research. The chapter concludes with an assessment of current applications and future prospects of research and work in AI.

2.1. Initial works in AI and the history of AI

While the notion of building “intelligent” machines can be traced back in philosophy, literature and the history of engineering for millennia, actual research on artificial intelligence began in earnest in the 1950s. In fact, the term “artificial intelligence” itself was coined in 1956, at the so-called “Dartmouth conference”, organized by John McCarthy, who is commonly recognized as the major founding father of AI.

This conference and the appealing ideas that were generated there sparked the development of machines and programs that were “intelligent”, trying to fulfill the agenda of AI in the years to come. This undertaking became mainly possible due to two technological breakthroughs that were made a decade earlier: The invention of the electronic computer and the storeable computer program.

Propelled by these groundbreaking inventions and by a sound theoretical foundation, which was provided by theorists like Shannon or Wiener, AI took an amazing pace in the late 1950s and early 1960s. The speed at which AI-research progressed gave rise to fantastic speculations and predictions about the revolution of daily life by AI. In fact, the whole field was very much excited by the fast successes in the past and assumed that AI would reach it’s goal of creating a truly “thinking” machine very soon, taking on a very optimistic outlook on the future.

The first concrete milestones in AI itself were achieved by Newell and Simon in the domain of problem solving. They developed a whole series of programs like “Logic theorist” or “General problem solver” that could solve well-defined logical problems (by making deductions) or being able to solve puzzles like the “tower of Hanoi” (by searching and operating on a “problem space”, the computational representation of the problem).

As these developments gained much public and scientific interest, government funding for AI research was readily available in huge amounts in the early 1960s - especially at MIT and peer institutions - further increasing the momentum of AI, leading to many innovative

programs and computer models in a wide variety of subjects. Most prominently are SHRDLU (spatial problem solving), SIR (language comprehension) or ELIZA (HCI and language), among many others.

As the direct impact of AI on psychology in the 1950s was minimal, the indirect impact of the developments in AI research was huge later on, reaching an apex in the late 1960s.

The AI-notion of understanding the human brain as an information processing device served as an important guiding heuristic for much of cognitive psychology. Ramifications of this notion lead to very intense cross-talk between cognitive psychologists and AI-researchers, elaborating on the details and constraints of this idea. This debate led to much conceptual developments down the road (scripts, schemata and mental models, to name just a few). Clearly, AI was truly a cognitive science in the 1960s and cognitive psychology derived many concepts from research on artificial intelligence. Most importantly, the emphasis on the importance of the “mental” representations to solve a cognitive problem that arises when dealing with the world stems from AI research, as many successful problem-solving programs (as most of those mentioned above) feature a “representation” of the problem space to perform their tasks. In parallel, more and more philosophers joined the debate, pondering the philosophical implications of concretely existing “intelligent” problem solving programs.

All of this changed in the 1970s, when more and more “knowledge-based” models were introduced in AI. In fact, knowledge-based models became the leading paradigm and the state of the art of AI in the 1970s and throughout the 1980s. This gradually led to three transformations: First, the claim of being able to construct useful “content-free” general purpose AI systems was dropped – the programs became more and more confined to be used in a circumscribed and limited domain, drawing from a large domain-specific knowledge-base. These so called “expert systems” like DENDRAL by Feigenbaum et al. (interpreting mass spectra of compounds in organic chemistry) or MYCIN (for medical diagnosis and treatment) were very successful in solving actual “real-world” problems, going far beyond

everything that AI was capable of until then. In fact, expert systems were so successful, that they became the bread and butter of future AI work for all decades to come. This led to the second trend, the development of programs that could solve problems, rather than the development of programs that would illustrate a point about human intelligence, largely abandoning the interdisciplinary nature of AI. Third, AI thereby literally made its way from the lab to the marketplace via expert-systems. By the 1980s, AI research was already almost exclusively commercialized, applied research for companies. The products generated by this research were adopted and used in a wide variety of fields: Serving automated quality control in production facilities, helping in medical practice and even the military (“smart bombs”), among others. This commercial trend effectively transformed the field from a cognitive science to engineering. The initial concepts and goals were practically dropped, while emphasizing the development of solutions for specific problems, using expert systems.

On the other hand, pure research in AI stalled to the present day.

This dramatic development has several reasons: Besides the huge amounts of money that were surely involved in drawing AI researchers into applied work, tempting them to give up on pure research, there were also a lot of problems in AI itself that could not be easily overcome, leading us to our next sections.

2.2. Problems of AI research

Given the ambitious outset of artificial intelligence and the fast pace that was maintained throughout the 1950s and 1960s, one might wonder, why artificial intelligence research would have taken the path outlined above. Obviously, the people in the field decided to go for the money, while dropping most or all philosophical and psychological issues that they initially also wanted to address on the way.

The reasons for this most unfortunate development are manifold:

First and foremost, the limitations of AI on several levels became apparent very soon:

Despite long-lasting, quick and steady improvements in the available hardware – sometimes improving performance on several orders of magnitude just in a couple of years – AI research could not keep pace with this momentum, just to the contrary – it stalled. To make matters worse, the problems seem to be not merely technical, but inherent in the very issues itself:

Initial successes in modeling symbolic manipulation tasks by Newell and Simon were prematurely considered as examples of programs that master the holy grail of human cognition: Logical and mathematical reasoning. Everyday problem solving was considered to be trivial in comparison. Hence the optimism of these early days. As it turned out, things are actually just to the contrary: The problems solvable by AI – namely symbolic manipulation tasks are trivial in their computational requirements, while the real hard problems are related to everyday actions in a complex and changing environment. The early pioneers in AI research totally underestimated the computational complexity of this kind of actions, rendering the early successes of AI research as irrelevant. The real problems are all those things that human subjects actually do with great ease, like generating a shopping-list and acting on it. At least, due the utter failure of AI to replicate this human performance, we are now – as psychologists – able to appreciate the computational quality of seemingly mundane problems like that.

Therefore, it became clear that humans operate on and use a vast database of background-information of interrelated world-knowledge to solve their problems and interpret new information. So far, no computer program was able to acquire the vast amounts of knowledge that an average human adult possesses and thus, all AI enterprises were sharply limited to domains of very constrained knowledge spaces. That's, how expert systems came about in the first place. They tried to overcome the problems that result from programs without a knowledge base, but were unable to provide a knowledge base that was large enough to do a wide variety of tasks. This led to the creation of programs that could solve domain-specific problems very well, but could only do that and nothing else. For example, even programs that outperform over 99.99% of human chess players and that draw from a large database of moves can't do anything else. Besides, they also don't go about their business in a human fashion, but mostly use "brute-force" approaches. Most programs in other domains to date are similar. This doesn't exactly map to our notion of "intelligence"; every janitor seems like an universal genius in comparison.

This raises another issue: While computers are obviously able to deal with well-defined problems that are limited in scope at an amazing speed, they are almost helpless in the face of ill-defined problems or problems that need a large knowledge-base or even involve a large amount of variance. Even though developments in "fuzzy logic" and neural network approaches in computer science, these problem remains unresolved: People are able to deal flexibly and appropriately with huge amounts of variance, computers are not – they are unflexible. This problem is well known to AI-engineers who try to solve even more mundane problems like the recognition of handwriting, voices or faces. People are so good at these problems that they don't even notice the problem, there is almost no cognitive effort and almost no errors. For AI, these are hard problems. Even if the respective computer-programs can do it after a lot of training, the error-rate generally remains very high, the programs are easy to trick by trivial manipulations or distortions and are still very confined to that specific

problem. As useful as these devices might be from an engineering point of view, it leaves us with nothing but admiration on how the human mind solves the problem of variance and flexibility at the same time – constantly adapting and acting in a goal-oriented way in a world full of variance, computational complexities and change. This particular issue still remains a total puzzle to contemporary AI researchers with a remaining interest in cognitive science – and they have largely given up on it – or so it seems. Even computers outfitted with the latest AI technology still lack any form of “common sense” to go about their business. This problem certainly won’t be solved in the near future, as all recent attempts to feed the computer with the necessary facts (via the internet) have fallen far short from the initial expectations.

To wrap this up, humans are able to handle multiple independent and unrelated problems, they are able to deal with variance, able to handle impoverished, truncated or distorted input and are generally flexible in their approach to solve problems and reach their goals. Computers still don’t do any of these things and most likely, this is due to a large difference in the amount of (implicit) knowledge about the world: AI only works if the “problem space” is confined to a rather small set of very well-defined items. Even though AI has become “knowledge-based” in a way, coming up with all kinds of expert systems in the 1970s, the knowledge base never proved to be sufficient to consider these systems to be “intelligent”. Far from it. As the insufficient knowledge-base is an obviously still one of the worst problems of AI, a real breakthrough in AI would consist in a breakthrough in machine learning. For some reason, computer programs proved to be surprisingly incapable of doing that on their own – even with the latest multi-layered neural network approaches. This accounts for much of their inflexibility as well as their limited knowledge base. A self-learning system could overcome all these problems if it is able to learn adaptively, while avoiding to learn the wrong things, like humans do. After all, superior human performance in everyday problem solving can largely be accounted for by continuous and adaptive learning, not only by “intelligence” per se. As this has even been already proposed by Turing himself (Turing 1950), this fact has

been largely disregarded so far, probably due to the decline of behaviorism and related theories of learning in the late 1950s, probably due to technical problems involving the fact that it's not easy for computers to evaluate the value of things that they learn (something that is not easy for humans as well). It's very easy to learn wrong and detrimental things. In any case, after over 50 years of research, progress is almost non-existent in this respect.

On top of these severe technical problems, there are even more profound conceptual and philosophical problems involved in AI.

At the heart of these issues lies the question: "What does it mean to be intelligent?"

If one wants to model human cognitive processes, he better has some hard criteria to assess the nature of these processes. As the controversy on the nature of human cognitive processes in general and intelligence in particular is still raging, it is unlikely that we will be able to come up with a consensually accepted definition what constitutes intelligence. Therefore, it's not entirely clear what is actually to be modeled. Is successful problem-solving performance sufficient to be considered "intelligent"? Most philosophers and psychologists would likely disagree. Related to this issue is the question if artificially created systems can ever give rise to true "intelligence" in a meaningful way or if they can only pretend to do so, simulating "intelligent" behavior, without generating it.

The Turing-Test (Turing, 1950) was devised to address issues like that. Turing states that if we can't differentiate between naturally intelligent and artificially intelligent systems, we should grant both systems the status of intelligence. The Turing-Test itself is an actual implementation of this idea in a conversational setting. However, Turings definition is merely formal and was sharply attacked by (among others) John Searle in his famous "Chinese Room" argument, in which Searle pointed out that symbol processing following formal rules (that's all computers do) can never lead to true "understanding of meaning" and intelligence, only to pretending of understanding and fake intelligence. The argument illustrates that even if the symbol-processing system is totally incapable of understanding any of the symbols in a

meaningful way, only relying on syntax in the process of generating an output, an outside observer who only sees the input and the generated output can't distinguish between a truly intelligent system that operates on the meaning of the symbols and a non-intelligent system that operates merely on the syntax of the symbols. The system will "look" intelligent to the observer in both conditions. Therefore, the Turing test is invalid to assess the intelligence of a system – non-discriminability between natural and artificial systems is not sufficient to constitute intelligence, as the artificial system can mimic certain characteristics of the natural system, seeming intelligent by doing that – without really being it. Even though the chinese-room argument itself has come under severe criticism, the issue is clear: It serves us as just one of many possible examples for the deeper philosophical problems that lurk within and surround AI research. All of these deep conceptual issues are far from being resolved. It even remains unclear, if they are solvable in principle. Facing this huge mountain of nagging problems, most researchers on AI have dropped the more ambitious goals of "strong AI" (generating true intelligence), humbly focusing on applied work to solve specific problems.

In conclusion, it became rather salient during the 1980s that we won't construct true "thinking machines" in the common understanding of the word (and as envisioned in the 1950s) anytime soon and if ever, not by the approach taken until now. For a more thorough critique of contemporary AI research, please see Newquist (1994).

2.3. Outlook and applications

Most of what is to be said in this section has already been implied in other sections, therefore this sections also provides a summary of our brief discussion of AI.

Due to the problems discussed in the previous section, the ambitious research of AI has indeed cooled down very much. AI has definitely changed it's focus from being a cognitive science that strives to build "humanly" intelligent machines and to understand the philosophical implications that derive from it, AI has taken a more pragmatic way in the recent decades, being almost exclusively an issue of engineering alone, getting the job done no matter how. This AI is largely "alien AI", not even pretending any more that the machine works the way humans do (or vice versa); the information processing in these models generally lacks any resemblance of human information processing, often using brute force approaches (chess programs are a prominent example).

Using the huge advances in hardware development, the focus is now on inventing tools that aid people in tasks that are either tedious or hard to do for humans, things that a computer can do much better, very much like robots for labor-intensive, mundane and dangerous tasks - like welding in the automobile industry.

This approach strives to help making life and work more efficient, comfortable and enjoyable, while giving up on the claim of being able to build a thinking machine, which would probably re-define our view on our own intelligence or what it means to be human.

Applications of this kind of contemporary AI can be found in many fields, including the entertainment industry (video-games), the military ("smart" bombs and drones), security (recognition of faces and voices), human-computer interfaces (recognition of handwriting, voices), language (translation) and many more.

As computer hardware continues to get more and more powerful, the AI-programs that perform these tasks will become better and better at it, while leaving the problematic issues of

the previous section largely unresolved, further supporting the idea that these issues are not only technical, but also very conceptual in nature.

Therefore, AI has failed to reach its scientific goals, it has succeeded in the solution of a large number of practical problems. More and more, AI is getting to be just another branch of technical engineering, without too much relevance for theoretical issues in contemporary cognitive psychology or cognitive neuroscience. There is not much interaction going on any more between AI and the cognitive sciences, not at any level.

In that way, AI has not fulfilled the fantastic promises and claims that were made in the 1950s. Moreover, these promises most likely won't ever be fulfilled in the pragmatic framework that actual research has now adopted; the agenda of AI has changed. Even most philosophers have withdrawn from the unresolved questions in AI. Without major conceptual breakthroughs and changes within AI itself, this state of affairs can persist almost indefinitely. If AI ever wants to model true human information processing, it has to go far beyond all the symbolic manipulation approaches that have been successfully invented in the past and probably build a structure of brain-like complexity in the process to do so.

A huge endeavor for the future and a long road ahead. Yet, any existing program is far from passing even the Turing-test and all of them are just pretending to understand a conversation, not really being able to do so.

What else can be said about the current status of AI and future prospects, especially in respect to issues of cognitive psychology?

In the very least, research on AI has re-defined our view of what constitutes a "hard" computational problem. While most cognitive theorists before the 1950s (and most people still today) considered mathematical and logical thinking as "hard" and the pinnacle of human intelligence, research on AI showed that mathematical and logical problems can in fact easily be solved by symbol-manipulating computer programs. It are the seemingly mundane tasks like perceiving, acting in a natural environment, performing motor acts and thinking about

arbitrary issues that are still surprisingly hard to do for any computer and robot. It has been particularly this research that made us appreciate the computational complexity of these problems and gives us an idea on how much computational resources our brain presumably devotes to these tasks to make them seem so easy and trivial for us.

The failure of the strong claims of artificial intelligence are very informative for cognitive psychology regarding our understanding of the nature of computational problems: The failure of AI to model certain performances can still enlighten cognitive psychology and cognitive neuroscience about the constraints and structural properties of the actual computational problems that humans are able to overcome in some way.

Moreover, the elaborate principles of AI developed so far can readily be used as tools in all kinds of other and unconventional undertakings, such as an understanding of creativity by working on things like “artificial creativity”. Artificial creativity in fact uses the methods of AI to understand creative information processing.

The whole rest of this paper will be devoted to these issues: Creativity and artificial creativity.

3 Creativity

3.1. Conceptual basics

As stated in the introduction, modern creativity research within psychology was basically started by Guilford (Guilford, 1950) who delivered a powerful presidential address to the APA on this topic in which he proposed more research on the neglected issue of creativity. His speech was successful, in that it actually sparked research on creativity, yet there are many unresolved conceptual problems concerning the nature and measurement of creativity.

Most psychologists would agree with the notions originally brought forward by Guilford himself (Guilford, 1967) that a creative act has to be **new** and **good/useful** in order to be called creative and that creativity differs from IQ and other problem-solving skills in that it is **divergent thinking** to solve a problem, in contrast to convergent thinking in problem solving.

Most contemporary textbooks commonly (see Galotti, 1998 or Kosslyn, 2000) share these definitions. Obviously, they are very compelling and plausible: Something that is only new, but not good or useful can easily be produced by anyone any time and is likely to be considered nonsense. Mere newness without value is just originality, not creativity. Something that is a good and valuable product, but not new is not really creative as well, just mere productivity or technical craftsmanship. Even worse, the infamous saying of

“re-inventing the wheel” illustrates that society doesn’t approve of undertaking like this. There is less agreement on almost everything else in the realm of creativity research; scientists construe the concept very differently. Most distinctively, creativity is an inherently interdisciplinary issue. Not only is creativity studied within different psychological fields like cognitive psychology, personality psychology and social psychology, but different academic fields like sociology, history and literature as well. On top of that, creativity doesn’t even remain an issue within academia as a whole. Compared with other constructs for example to IQ, where psychologists can define as intelligent what they please, the power to decide what constitutes a creative product rests – by definition – not only with psychologists or scientists,

but with the gate-keepers of the respective domain: Art critics, other authors, other composers or – in some cases – even the general public at large.

As creativity is therefore a highly complex phenomena, we can distinguish at least 4 aspects that are quite different and are usually studied differently within different disciplines or with different methods within the same discipline:

-The creative **person**. The big question here is, what makes a person creative and why are the differences in creativity between persons so huge. Questions like that are normally raised and studied within personality psychology or even clinical psychology. Commonly used methods are the analysis of case-studies, the use of questionnaires, creativity-tests and the like.

-The creative **process**. Here is the big question, which cognitive processes lead to creations that are considered to be creative. Cognitive psychologists take interest in these matters and generally try to use an experimental approach to address their questions.

-The creative **product**. What makes a product creative? What are creative products? Are specific products creative and how? Questions like this are pondered by sociologists, the historical sciences and the gate-keepers of the respective domain - those that have the power to declare a product to be creative if they deem it to be creative.

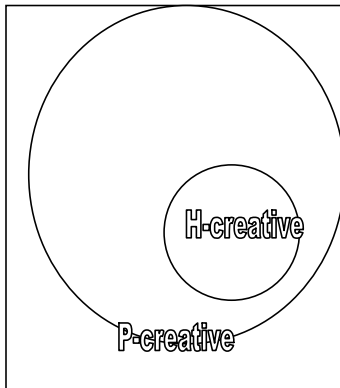
-The creative **environment**. What environment fosters the occurrence of creative events? How can creativity be facilitated? Social psychologists and school psychologists and educational researchers primarily deal with this issue, using a wide variety of methods.

Evidently, creativity research, the strive to get a comprehensive understanding of creative thinking and creativity is an inherently interdisciplinary enterprise. If someone constrains his research or thinking to a certain field, he will most likely only be able to answer questions within this field, for example questions concerning creative processes. As these “smaller” questions are hard enough themselves and more likely to get answers than the big (and mostly meaningless) “what is the nature of creativity” like a philosopher would ask it, this is fine, but should be kept in mind.

On top of that, I want to point out another important conceptual distinction that was brought forward by Margaret Boden (see Boden, 1992). As creativity is constituted by products that are new and useful, produced in the creative process by the creative person in the creative environment, one can make several interesting differentiations: What is considered “useful” (of high quality) in a certain domain differs from domain to domain, it can be the solution to a problem (in science) or something beautiful (in art) or something else and rests largely with the gatekeepers of the domain and has to be studied idiographically in every instance. The quality and value of an idea can't be assessed within the science of creativity itself – it relies on science-external factors. On the other hand the novelty of something being “new” is also not unambiguous. Boden differentiates between historically new (**H-creative**) and personally new (**P-creative**).

H-creative products are new to everyone, to humanity as a whole and are generally considered as “truly creative” by the common public. Historians and sociologists of science are only concerned with this type of creativity. P-creative products on the other hand, are not new to everyone, only new to the person itself. This “everyday-creativity” isn't considered creative at all by most people, but if one is interested in the creative process, it shouldn't matter if someone else also had the idea before, if the “re-invention” is really made independent of everyone else. Therefore, cognitive psychologists are most interested in this kind of creativity: It also fits the definition of creativity (new and useful) and is far easier to study in the laboratory. Moreover, it's very hard to assess, if H-creativity is given in a particular instance: Sometimes, an idea is considered as H-creative but is not (someone had the idea before, but this wisdom was lost), in other cases something is H-creative but denied that status by the gate-keepers. Please note that even P-creativity is not always easy to determine: Everything is new to everything before in most (but most likely not all) aspects. Of course, new ideas derive from personal experience, often reconfiguring or re-structuring things that are old, not new. Therefore, quantifying the novelty of an idea in a coherent way is not as easy and compelling

as it seems at first glance. In any case, H-creativity and P-creativity are in an asymmetric



relationship. Everything that is H-creative is also P-creative, but not vice versa, as illustrated in figure 1.

In the everyday business of science, this distinction shouldn't matter, but if certain fields and findings are in conflict, it's important to keep in mind that this distinction has implications on how to resolve the conflict: Maybe these fields study different

Figure 1

things in the first place and there is necessarily nothing

inherently odd or contradictory in the nature of creativity itself that needs to be resolved, or nothing bad about a particular way to study creativity as these conflicting findings are artifacts produced by the ways in which we approach the study of creativity.

This approach allows multiple "truths" about different – but closely related – objects of study. Seemingly contradictory findings are not necessarily contradictory in this view, but can be accounted for by different methods of study, different operationalizations and a different emphasis on certain aspects of creativity. This caveat should provide a successful defense against radical relativism and also provide insight in how to resolve some of the many unnecessary struggles that are so typical in creativity research.

After all, the definition (new and valuable ideas) of creativity as brought forward in this section is very general and very formal. It is certainly not the only possible definition and it's freeness of specific content and independence of a specified physiological basis allows for a very wide range of operationalizations and applications as well as certain ways to study creativity, like in the computer models of creativity that we will see later.

3.2. Theories of creativity

The cautious approach to clarify conceptual issues before going into the subject in detail taken in the last section is really necessary in the case of creativity, as shall be seen in this section. It is necessary to avoid all the troubles that immediately ensue unless the conceptual basis is solid. This is true for every scientific undertaking, but even more so for the study of creativity. To simplify matters, I will start by focusing on the aspect that is of most interest to cognitive psychology: The creative process. I will therefore purposely shorten the discussion of other aspects or viewpoints that deny the relevance of the creative process in the first place like Sternberg, who conceptualizes creativity as a skill (see Sternberg, 1988) or Simonton who basically proposes that creativity is an ill-framed epiphenomenon of productivity (see Simonton, 1988 or Simonton, 1997). These approaches are undoubtedly very interesting and the arguments about other aspects of creativity are no less fierce, but for the sake of brevity, we have to skip most of what is not relevant for a cognitive psychological understanding of the matter in order to advance to the computer models of creativity.

3.2.1. The creative process – early models

Guilford (1967) himself proposed the following basic 4-step creative process (Figure 2):


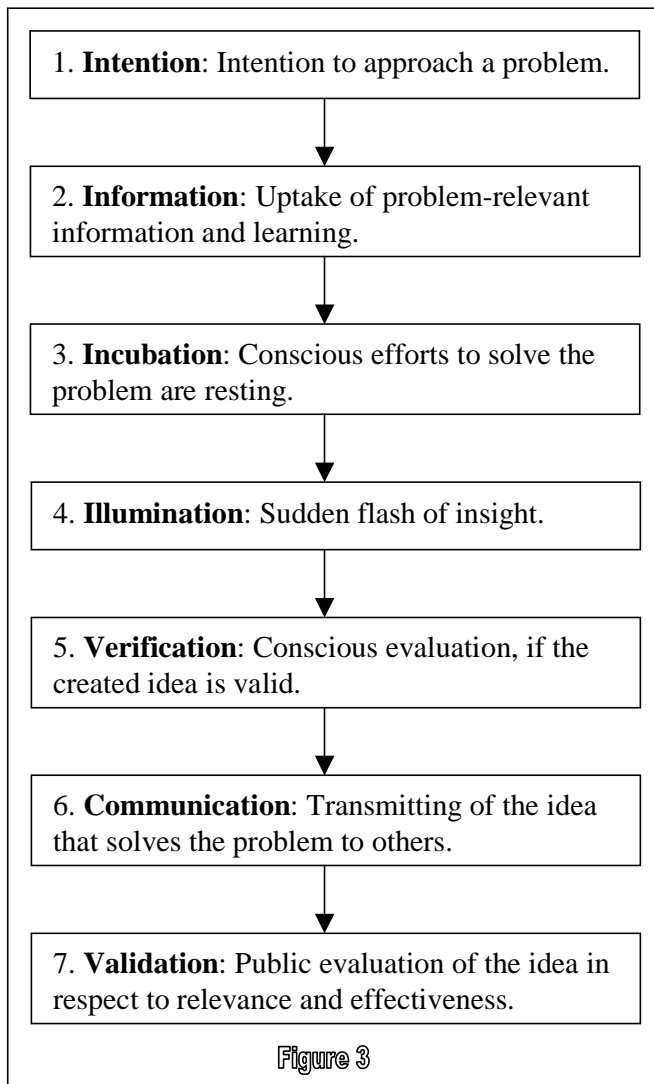


Figure 2

The shortcomings of this approach are obvious: Even though the model provides a first heuristic to think about creativity, it's far too unspecific for concrete research on creativity, too untestable and merely descriptive on a very abstract level. It doesn't say, how to progress from one step to the next, how the steps come about, about the time between steps or explain anything at all.

Therefore, a more elaborate model was adopted by most creativity researchers to guide much research. This model is also a multi-stage model and is a modified version of the model proposed early on by Wallas (see Wallas, 1926). This model was so successful because many research traditions are represented within it, it explicates intuitions about creativity and scientists with different interests can place their emphasis on different stages of the model. For example, Feldman (Feldman, 1988) emphasizes the Incubation and Illumination-stage. Csikszentmihaly (1996) emphasizes the importance of the communication and validation



stage. The model is exemplified in figure 3.

While there is some empirical support for the notion of these stages, (see Metcalf & Wiebe, 1987) this model has recently come under severe critique from different factions: Weisberg (1993) and Perkins (1981) criticize that it's relying on mysterious processes like incubation or illumination. Others, like Csikszentmihaly (1996) point out, that it biases our thinking about creativity by implying a largely cognitive, largely static and serial undertaking while underestimating

social processes. These discontent people often proposed alternative models, some of the more fashionable ones we explore now.

3.2.2. Alternative models of creativity and the creative process

There are many other descriptive models of creativity that are either simpler (and therefore generally better empirically testable) or even more sophisticated. Most of them abandon the notion of incubative or illuminative stages or processes, considering them as rather esoteric. Note that this illumination (the A-ha experience) is not part of the formal psychological definition of creativity, yet an integral if not constitutive part of the creative act in folk-psychological conceptions of creativity.

A prominent version of a very simple, reductionist model is that of Campbell (Campbell, 1960). Campbell proposes only two stages:

1. Generation of as many different possible solutions to the problem as possible, without evaluating, censoring or rejecting them.
2. Interpretation and evaluation of these possibilities in respect to the problem at hand.

This model is employed in techniques to increase the output of creative ideas like brainstorming (Osborn, 1953) or as a general technique in problem solving (“create and test”).

While this approach is merely descriptive, Finke (Finke et al., 1992) took this model to a new level, explaining creative achievements within his “**creative cognition**” framework. In this cognitive psychological framework, he merely relies on standard cognitive processes like memory and imagery to account for creative performances. His models builds on the adaptive alternation between divergent and convergent thinking in order to solve complex problems. Divergent thinking alone doesn’t solve problems in a coherent way, as it tends to increase the number of ideas, not the quality and convergent thinking alone is not creative as it increases the quality, but decreases the number of ideas.

Other scientists in the field of creativity suggest different models that go beyond cognitive psychological accounts of creativity and merely cognitive processes to explain why some people’s work is considered creative, while others isn’t. Prominently, Sternberg & Lubart (Sternberg, 1999) propose an **investment theory of creativity**. This theory compares the

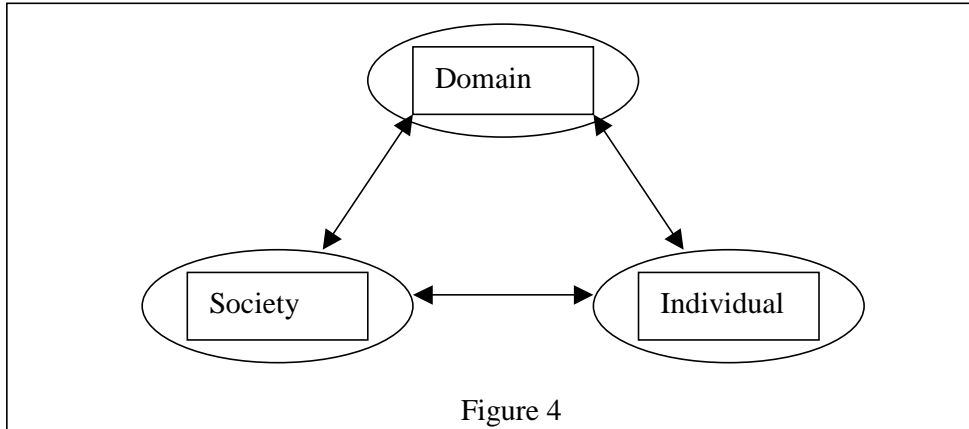
creative person to an investor, emphasizing that the whole configuration of mental skills of this person and the use of these skills in the interaction with the environment gives rise to creativity, no specific cognitive processes per se.

Within this framework, creative people are described as tending to take risks, by putting effort in neglected and underrated ideas. Then, they try to persuade others of the value of these ideas (acting like an investor who acts by the motto “buy low, sell high”, hence the name of the theory). Therefore, a creative person has to resist the opinion of the majority first, see the potential of an idea and be persistent enough to convince others of the value and validity of his notions. The view of Lubart & Sternberg, features different creative “resources”. Resources that lead to creative performance. They differentiate 6 types of resources: Cognitive resources like knowledge and intelligence, affective resources like personality and motivation and additional resources like cognitive style and the environment. Interactions between these resources give rises to creative skills. This perspective integrates many empirical findings about creativity, namely that a high intelligence, knowledge about a domain, a persistent and risk-tasking personality and a intrinsic motivation is improving creative performance. A “legislative” cognitive style that tries to achieve general laws and a supportive environment that approves creative works are also fostering creativity.

The appeal of this theory is indeed that it integrates many empirical findings, while accounting for the complex phenomenology of creativity and explaining, why only so few people are highly creative (because it takes the perfect combination of all these resources). The drawback of this theory is the conceptual fuzziness, leaving many parts empirically untestable and unfalsifiable.

While Finkes model can be described as a paradigmatical case of a cognitive psychological model and Lubart & Sternbergs model as a typical model of personality psychology, there are prominent models that go beyond that, for example the “**systems view**” of creativity, proposed by Csikszentmihaly, emphasizing the importance of social skills and social

processes in the generation of creative acts. He tries to explain, why some cultures are far more creative than others, essentially considering creativity as “cultural evolution”. Moreover, this approach doesn’t view creativity as an intra-psychic phenomenon, but rather as a complex relation between an individual, the domain and the society (see Figure 4).



The domain is a system of representation of wisdom and knowledge with certain rules and contents. The society (or “field) consists of a group of gatekeepers deciding, what belongs to the domain and what not.

The interactions in this system are rather complex: The individual learns about the methods and issues in a domain and generates novel ideas, trying to change the domain. These novel ideas are evaluated by society. If approved, they become a part of the domain and are passed on to future generations of individuals. As pointed out above, this process bears close resemblance to evolution and can be understood as cultural evolution.

Creative products are those changes in the domain by the individual that are retained by society. This model allows for many possibilities to influence the probability of creative acts: A culture can restrict the access to a domain, a better notation system in the domain fosters creativity (e.g. arabic numerals vs. roman numerals), a society can be more or less hierarchical, etc.

The appeal of this model is, that it can account for the large inter-cultural differences in creative achievements and that it is consistent with most empirical facts about creativity (like

the fact that people have to learn about the ways of a domain, before they can significantly contribute to it), making sense of these empirical findings. The drawback is once again the lack of concreteness and empirical testability.

3.3. Empirical issues

In this section, I will briefly review solid and interesting empirical findings that were achieved by creativity research to date. These rather arbitrary findings can be reconciled and understood in light of the theoretical frameworks discussed above.

There are all kinds of interesting and highly relevant empirical questions in the field of creativity research. These are important for individuals as well as for society as a whole. For example: Can creativity be increased or fostered? If yes: How so? Is creativity an inherent trait? What is the relationship between other psychological constructs like IQ or other psychological conditions like schizophrenia and depression and creativity? Can everyone be creative? Is there a specific time-course for creative products? And many, many more.

Still, more questions remain concerning the nature of the creative process itself: What is the role of experience, motivation and rather esoteric processes like illumination and insight in creativity, what the role of unconscious processing?

As interesting as these questions might be: It's extremely hard to address or resolve these questions on empirical grounds, if the theoretical and conceptual foundations are so weak and contested as in the case of creativity. Without a coherent theory of creativity, the measurement of creativity is very hard. Without a reliable and valid method of measuring creativity, the chances of resolving disputes by empirical facts are slim.

In any case, this brief review is – by necessity – extremely selective and only reviews the most important or most solid evidence. It's important to keep these caveats in mind throughout the later parts of this section.

3.3.1. Methods to study creativity

As pointed out above, measuring creativity is one of the harder problems for psychological methodology. Currently, there are two very different and widely-used approaches:

One of them is the analysis of **biographies**. These approaches are genuinely idiographic, studying one or more individuals in great detail. Authors working in this tradition then generalize their findings from this case-studies to claims about creativity in general. The positive thing about this approach is its construct-validity. As only the biographies of people that were undoubtedly creative like Beethoven, Einstein, Darwin and the like, this approach can avoid the problem of assessing the quality and value of the creative ideas in a scientific way. The biographies are analyzed in terms of personal traits of the creative person, the background, environment, life events and their impact on the creative performance of the respective individual.

As most of this research is carried out after the death of the creative individual, research in this tradition focuses on the creative person and the creative environment.

Problems of this approach are manifold and obvious, it inherits basically all the problems that are entailed by case-studies: First of all, case studies are always anecdotal – it is entirely unclear, if the selected individuals are representative of the population of creative individuals, leaving room for biased interpretations. Moreover, these approaches tend to be very qualitative, interpreting complex patterns in hindsight. As the variance in these patterns is generally very high and the patterns are generally very complex – creative individuals don't necessarily share any traits or backgrounds – this approach is not too informative for a psychological understanding of creativity, as it is hard to draw general conclusions from these case-studies. Of course, studying the creative process or establishing causality is also impossible within this framework.

Simonton tried to overcome some of these problems within his “**historiometric**” approach to creativity, as he analyses biographies of many creative individuals in respect to formal and quantifiable parameters, trying to establish formal laws of creativity (Simonton, 1997).

As the biographical approach suffers from this many inherent drawbacks, a complementary methodological approach is very popular in the psychological study of creativity:

This conflicting approach is test-oriented and grounded in the **psychometric** tradition. The advantages of this approach are, that one can test “normal” people and quantify their creative performance or ability to be creative, allowing for experimental studies and the study of the creative process itself. The main problem of the test-based approach is construct-validity.

It remains an open question, if so-called “creativity tests” (in analogy to IQ tests) actually measure “creativity”. There are virtually no studies that correlate test performance with actual outstanding creative achievements in “real life” in a systematic way. High test-scores in these tests don’t necessarily lead to creative achievement in life and individuals that are commonly regarded as creative generally refuse to take part in these tests.

Moreover, most creativity tests, prominently the Torrance test of creative thinking (Torrance, 1974) just measure verbal productivity and divergent thinking, neglecting convergent thinking, evaluation of the quality of the produced ideas and other issues that are most likely to be involved in creative information processing. Many tests just feature an index of how many ideas one can come up with after given a verbal instruction to do so. The Torrance Tests of creative thinking include 6 verbal and 3 figural subtests with tasks like “alternative uses” or “figure completion” where subjects have limited time to come up with as many solutions as possible. Under these conditions, it can be doubted, if these tests actually assess processing that plays a role in the generation of commonly recognized creative products. In addition, social constructivists would argue their concern that these tests can measure – at best – resources for creativity, not creativity itself, as creativity is a socially defined construct, only existing as a complex social relation.

Therefore, both approaches embrace major problems – suggesting to take the empirical findings that were obtained using these approaches with a grain of salt.

3.3.2. Motivation and creativity

The major player in the investigation of the relationship between motivation and creativity is Amabile. The main finding so far Amabile (1985) is, that “intrinsic motivation” fosters creativity, while extrinsic motivations disrupts and suppresses it. This is explained in terms of productivity: Individuals with an intrinsic motivation are more persistent in their work, as they are continuously driven by their “inner” urge, while all external and extrinsic motivations are far more unstable. Moreover, extrinsically motivated individuals just want to get the job done as soon and directly as possible, shortening the divergent thinking processes of creativity. If extrinsically motivated, the individual tends to evaluate ideas earlier and look if they solve his problem. They are also more likely to accept the first solution that comes to their mind. Both trends hamper the generation of new ideas and therefore creativity itself.

3.3.3. Productivity and creativity

The investigation of productivity and creativity is intimately linked to Simonton. Given the definition of creativity, mere productivity is a huge confound. Simonton devoted his life-work to resolve problems that arise from this confound. In his paramount work (Simonton, 1997), he analyses the output of many creative individuals in many different domains (sciences, music, literature, etc.) over their lifetime. While assuming that the quality of each work is constant (because of lacking criteria to assess that), he finds that the amount of published output over a lifetime is not. By fitting curves to the data he got, he could even derive formula of expected future creativity of an individual, based on past productivity, the current age and the respective domain – individuals in literature tend to be more creative later in life, in contrast to mathematicians or physicists who achieve their major work early in life.

Moreover, he found that creativity is intimately linked to productivity. Individuals produce their most renowned works, when they also produce their most work. The individuals that are most regarded for their creative achievements in a domain are also the most productive ones. Simonton shows that 80% of total publications in a field are created by only 20% of the people in the field. These people are also likely to be considered as “creative”. Simonton explains this by stating that the individual himself is not capable of predicting, which of his works will be considered to be creative. The best strategy under this conditions would be to put out as much work as possible, to optimize the chances that some of the works will be regarded as creative by peers.

Of course, this model is only descriptive – mathematical modeling of life-courses can only lead to descriptive models. Therefore, we don’t know what causes these very interesting trends in the output of creative achievement-curves of highly creative individuals.

3.3.4. Knowledge and creativity

Most people would consider knowledge and creativity to be opposites. They think that the creative individual has to guard himself against knowing too much – knowing too much constrains the thinking along narrow lines, alienates the person and dooms creativity.

This view is still very popular among people who want to believe that and was supported for a long time by studies in the tradition of Luchins (Luchins, 1942), showing that “mental set” can prevent people from recognizing the best solutions. And as every thorough education can be conceptualized as establishing “mental set”, it hampers creativity... or so it seemed. Weisberg (Weisberg, 1986, 1993) deserves much credit in destroying this popular myth. In his many books and articles, he shows in detail, how experience and knowledge actually fosters creativity. In all cases, experience and intimate knowledge of a domain was positive, if not necessary for outstanding creative achievements. This holds not only true for creativity in science, but other fields as well, including music.

As a summary of his lifetime-work on the relationship between creativity and knowledge, he postulates the “10-year-rule”: Every individual has to work hard in a domain for at least 10 years before being able to contribute to the domain with a major breakthrough. He showed that this rule also holds true for so-called “geniuses” like Einstein or even “prodigies” like Mozart. The secret of these geniuses is – according to Einstein – their ability and willingness to invest 10 years of deliberate and hard work, practicing their domain-relevant skills every day of these 10 “silent” years, learning and mastering the methods and ways of their domain. It’s unlikely that this – rather optimistic view (everyone can be a genius!) can account for all complex effects that we observe in creativity, but his 10-year-rule is a very fine piece of scholarship and certainly helps to increase a psychological understanding of creativity. This findings can even be reconciled with the research by Luchins: Of course, creativity can be conceptualized as breaking mental set or overcoming functional fixedness – especially in the concrete creative process itself, but this doesn’t neglect the benefits of a decent education.

3.3.5. Personality and creativity

Obviously, personality traits that make hard, deliberate work and intrinsic motivation more likely make creativity more likely. This can be learned from the first parts of this section. But there is more to the creativity personality than that.

Sulloway (1997) convincingly showed, that a “rebellious personality” is very important for creative breakthroughs. There are many ways to develop a rebellious personality, most prominently to be “later-born” or discriminated by society. He points out, that all major revolutions in science and politics (including Darwin and the french revolution) were spearheaded and supported by later-borns, while first-borns are more likely to defend the established ways of the domain and the society. Sulloway provides a very elaborate argument to defend his interesting hypothesis. However, he can’t really explain, why so many later-borns don’t become rebels or creative. Basically, he capitalizes on rather weak correlations,

using huge sample sizes. Therefore, this approach won't account for much variance in creativity – the effect is rather small.

An effect that is in fact rather strong in the relation of creativity and personality is IQ. Contrary to popular belief, IQ is very important for creative achievement. An IQ under 120 almost perfectly correlates with creativity. Over the threshold of IQ 120, there is a null correlation between IQ and creative achievement. Obviously, an IQ of 120 is sufficient to do everything one wants to do and creative achievement is determined (and constrained) by other factors. Under the threshold of 120, IQ is apparently one of the main constraints of creativity.

3.3.6. Biological and clinical issues

Unfortunately, this aspect is still open to much speculation and popular myths. Biology can't tell much about the causes of creativity. There is some speculation about the relevance of the right cortical hemisphere for creative thought or the "integration" of both hemispheres for creative achievement, but any convincing empirical proof of this unspecific hypothesis is still lacking. Moreover, while there is some suggestive evidence for genetic factors in creativity – illustrated by highly creative families like the Bernoullis – we still lack any clue on which specific genes could be involved in creativity.

Clearly, biology could provide much more insight into the causes of creative achievements, without doing so yet. Personally, I would expect permanently elevated levels of neurotransmitters like ACh in highly creative individuals.

Concerning clinical issues, mental illnesses – especially manic depression and schizophrenia have been linked to creativity, as these illnesses can generate patterns of divergent and convergent thinking, that are presumably aid creative thought.

However, empirical evidence for such a relation is very weak, mainly confined to case studies and anecdotes of mentally ill, yet highly creative individuals (like van Gogh).

Biological and clinical aspects of creativity remain a mystery – the potential of these fields to increase a scientific understanding of creativity in the future is huge.

3.4. Summary

A scientific understanding of creativity is without any doubt extremely important for the welfare of our society and every single individual as well. It could be the key that opens the gate for a new, golden era of humanity. For example, an understanding of scientific creativity, on how new discoveries are made could revolutionize the way, in which we do science itself. Yet, it remains an extremely tough field of investigation. This state of affairs is very similar to many other issues in cognitive psychology like attention, categorization, reasoning or problem solving. These fields are equally important for humanity and also very hard to study.

Concerning creativity, many important issues remain unresolved to the present day. The picture drawn in section 3.3. is far too rosy – in fact, those issues are not so clear-cut.

On conceptual grounds, at least three different views on creativity within psychology are still alive and well – competing with each other, namely the approach emphasizing unconscious processing and insight, the cognitive psychological approach stressing “mundane” cognitive processes and social psychological views that conceptualize creativity as an emergent social construct.

Based on the available evidence, it would be premature to decide on the validity of these theoretical frameworks. This decision has to be postponed until solid empirical evidence allows to rule out some of the alternatives.

Please note again, that the conventional conception of creativity (new and useful) remains merely formal. It can be applied to almost any human endeavor, virtually any field of achievement or performance, including dance, art, science, literature, thought, problem solving, etc. Creative acts within these domains might share no single specific feature, could be radically different instantiations and manifestations of this formal definition. This means,

that a unified theory of creativity is not necessarily possible, as creativity might be an emergent property, working differently in different domains and at different times.

The everyday understanding of creativity as an unitary concept doesn't necessarily map on a scientific understanding of creativity – as science tries to understand and model what is actually going on in the world, dealing with very complex systems and huge amounts of variance in the case of creativity.

Therefore, these question remain unresolvable, so far. At least, not with the empirical data at hand. The problem is, that the methodological arsenal to study creativity is extremely poor and problem-laden in the case of creativity research, as pointed out earlier. Without adequate methods, the necessary empirical data won't be gathered. To make matters worse, one promising field of investigation for new research in creativity hasn't do anything at all, yet. A (cognitive) neuroscience of creativity is non-existent to the present day.

Prospects are better in respect to another, complementary science that could lend methods to creativity research. In fact, computer models of creativity promise to bring salvation to the troubled field of creativity research. With a methodological arsenal that is so poor, researchers on creativity can indeed use every help that they can get. And instead of abusing the powerful methods that were developed within AI for engineering alone, they could as well be used for a better scientific understanding of creative processes.

That's, what the next chapter will be all about:: „Artificial creativity“, computer models of creativity. After all, a psychological understanding of creativity so important that we can't miss any chance that could possibly advance the field beyond the sorry state that it is in now.

4 Artificial creativity

4.1. Foundations and conceptions

Artificial creativity is the attempt to apply principles of AI to creativity by using computer models of creativity. One of the most amazing things about computer models of creativity is, that they exist in the first place. Commonly, computers are viewed as the total opposite of everything that is creative: They are seen as the impersonation of everything that is mechanistic and deterministic.

The very existence of working computer models of creativity is an existence proof that creativity can in fact be considered as a computational problem. But even if one considers the brain as a biochemical computer and creativity as the computational problem, it's puzzling how creativity is possible – how the computational problem can be solved: The tuning properties of most neurons are fixed: How can this system come up with something new and evaluate the new product in an adaptive and meaningful way?

Specific computer models of creativity offer the chance to make inferences about how human creativity might work: Both systems, human and artificial face the same problem in the generation of creative products; However, our theories about how human creativity works suffer most from being fuzzy and empirically untestable. This is not the case for computer models of creativity: They have to be extremely concrete and specific in order to work, complementing models of human creativity and allowing inferences about how human creativity could work.

In order to understand computer models of creativity, another conceptual distinction between different kinds of creativity becomes necessary. Boden (Boden, 1992) distinguishes the following possibilities to be creative:

1. **Combinatoric creativity.** This refers to combinations of ideas that are already known to the individual. This kind of creativity is important in associations and the forming of analogies. In computer models, this type creativity is often implemented as neural networks.

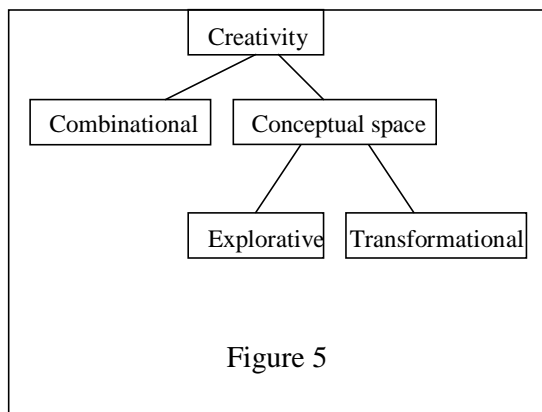
These networks are able to reconstruct a familiar object if given only parts of it or able to note similarities between different objects, allowing for the ability to draw analogies. However, these models are not yet really sophisticated and barely resembling human creativity. The lion share of creativity rests with the programmer, as he builds the neural net and evaluates it's output. The neural network itself is incapable of creating new nodes within the network or understanding, what it is doing.

2. Explorative-transformational creativity.

These types of creativity take place within a conceptual space. Conceptual spaces are established by constraints that structure the set of ideas that are principally possible. Different domains have different conceptual spaces generated by different sets of constraints.

This allows for two kinds of creativity: **Explorative** creativity consists in a search of the conceptual space, finding or realizing some of the possible ideas. **Transformational**

creativity alters the constraints themselves, changing the conceptual space and making new



ideas possible that are unthinkable of with a different set of constraints. Humans can do both types of creativity. Computer models are most often restricted to explorative creativity.

If one allows the computer to alter it's conceptual space, the results are often not

meaningful or valuable anymore, as the computer has no valid criteria for the value of ideas.

4.2. Specific examples of computer models of creativity

As creativity can be found in almost any human endeavor, this is also the case for computer models of creativity. Some prominent examples of such models are reviewed in this section.

All the examples have been taken from Margaret Boden (Boden 1992, 1996, 1999 and 2000), who is one to the leading scientists when it comes to computer models of creativity.

4.2.1. Painting and drawing

This field uses mostly models of explorative creativity and is very successful at doing so.

One of the most prominent examples is “AARON” by Cohen (Cohen, 1981).

AARON specializes in line-drawings (see figure 6 for an example). While every single

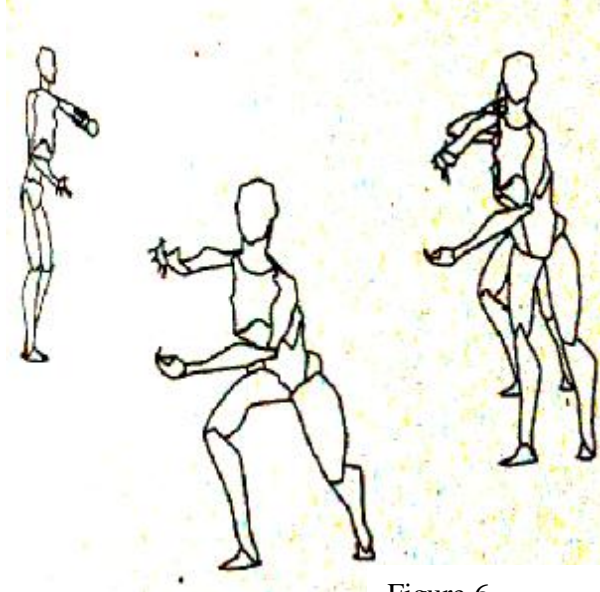


Figure 6

instance of AARON is confined to its conceptual space, it can produce an infinite number of different, but similar drawings in the same style – with great success: Drawings by AARON were shown in acclaimed exhibitions and were actually sold at art-auctions. A similar performance by human creators would undoubtedly be considered as a “creative”

achievement. Moreover, naïve subjects are unable to judge whether AARON-drawings were produced by a program or a human artist – thereby passing an analogous version of the




Figure 7

Turing-test (for creativity). Advanced versions of AARON are able to draw colored and more sophisticated paintings (see figure 7).

While the performance by these programs is astonishing, there are remaining problems:

First of all, those programs never change or improve their style, as they are confined to a constrained problem-space that they are merely exploring in a more or less random fashion.

These programs lack criteria to evaluate their own work and to change the style accordingly. External judges (humans) evaluate the output of

AARON and its siblings, forcedly changing the style to improve this output. Changes in style are nothing that the program can do on its own in any productive manner: Trials to make AARON more flexible and giving it the power to change its own conceptual space proved to be disastrous and quickly resulted in drawings without any coherence at all.

Therefore, these program remain bound to their conceptual space – confined to line-drawings of humans infinitely, doing anything else is totally out of scope. This doesn't seem to be too creative, as the program can never leave these narrow boundaries, forever drawing variations of the same theme, even after millions of rehearsals. The program itself doesn't even "know" about its creative performance – it bases on a random exploration of the given conceptual space. The attribution of creativity relies on and their **interpretation** of this output. The program itself doesn't try to express itself, follow any goals or perform self-reflexions – something that any human artist would undoubtedly do. Thus, even this seemingly very successful domain for computer models of creativity remains awkward in a way.

A breakthrough to meaningful transformational creativity (something a truly creative human artist does) relies on the ability to provide a value system for the program. If such a system would exist, the program would be able to change and improve its style in a meaningful way. However, as even humans have a hard time coming up with a value system for the evaluation of art, this is unlikely to happen. At least not in the domain of art and drawing.

4.2.2. Music

As music is a fairly constrained conceptual space – with a very limited number of musical notes, instruments and rhythms it is not surprising, that music is one of the most successful domains of artificial creativity. There are now programs that can generate and compose music in a wide variety of styles – even Jazz is not out of the reach.

Generally, these programs analyze rhythms and melodies of existing works in the respective style and generate new works by more or less random variations of recombinations.

Contemporary compositions by artificially creative music programs are really sophisticated and virtually indistinguishable from human creations for non-professional ears. There have already been cases, where professional musicians have incorporated parts of artificial creations in their repertoire. Unfortunately, I can't provide hearing-samples of these amazing programs in the framework of this paper.

As all existing programs that create music base on explorative creativity the same restrictions as in artificially creative drawing or painting apply. However, as transformational creativity is very rare even for human creativity in music, this fact is not as obvious as in the visual arts.

4.2.3. Science

Creative programs in the realm of science share much resemblance to expert systems, as they fully and systematically explore a small and very constrained conceptual space. An example of one of these programs is META-DENDRAL which is endowed with a conceptual space that includes physical-chemical constraints of molecule configurations. The task for the program is to look for new possible molecular structures. In fact, META-DENDRAL came up with configurations that were new and useful for humanity, therefore even reaching H-creativity. Thus, META-DENDRAL was cited as co-author in several papers in peer-reviewed journals. Yet, this undeniable creative success derives from a merely explorative creativity approach, thoroughly exploring the conceptual space.

Later versions of META-DENDRAL were even able to propose experiments in organic chemistry to fill gaps in the literature and to hypothesize about possible outcomes. Some of these proposed experiments were actually carried out with great success.

Other programs like "BACON" are able to analyse experimental data on it's own. BACON is outfitted with a database of mathematical and physical rules and is able to look for patterns in the data, looking for more and more complex relationships, as no simple patterns in the data can be found – very much like a human scientist would do it.

It can be contested, if these programs (and their many cousins) are truly creative, or if they only foster human creativity in a tool-like manner. The creativity of human scientists often consists in the generation and transformation of the conceptual space and the ability to raise new questions, not only to search a given conceptual space.

None of the creative programs to date is able to do that. Therefore, these programs can free the scientist from the more mundane aspects of creativity in science: Exploring all possible relationships between existing facts or analyzing data, allowing him to focus on the “higher” tasks of transforming his field.

So far, the programmer has to tell these programs, what the subject, the question and possible data are. The programs don't come up with that on their own. This means, that these programs are limited to the simulation of specific aspects of creativity, which can be very useful in their application in the natural sciences.

Besides these approaches, there have been attempts at computer-generated transformational creativity. It is no coincidence that this attempt has been made in the domain of mathematics, where the main problem – establishing a computable criterion for the usefulness of an idea – is easily solved: The rules of logic provide coherent and valid criteria for this. In addition to that, the rule to chose the simplest solution helps to decide between different logically valid options. As these programs try to simulate transformational creativity, they are able to modify themselves in crucial parts of their code, using generic algorithms.

Generic algorithms work by using an evolutionary approach: The program starts with a base and applies random changes. The results are then evaluated in respect to the criteria and value systems that are implemented in the program. The “best” mutation serves as the parent program for the next generation of random changes. After many iterations, the program is able to solve specific problems, the modifications change the conceptual space of the program, allowing it to do things that were not possible to former generations of the program.

One of these programs is “AM”, it builds on a database of 100 simple mathematical axioms and a system of criteria to evaluate new ideas, searching for new mathematical concepts.

In fact, AM was able to come up with many new mathematical concepts like addition, multiplication or prime numbers. Please note, that the program independently discovered or invented these programs and has currently rehearsed the first thousand years of mathematical history on it’s own and is still working on new theorems.

Again: The tremendous creative success of this program in mathematics is not that surprising. After all, if mathematics is the purely logical space it claims to be, there is no need for human mathematicians in the first place. Mathematics is a field that can potentially really be totally mastered by artificially creative computers. This is owed to the special structure of mathematics: All possible “discoveries” and theorems are – by definition – already implied in the axioms.

The only other field where transformational creativity could be promising is also one with hard criteria for the feasibility and value of an idea, namely engineering. One of the programs “EURISKO” was able to invent a three-dimensional chip, which was promptly patented.

Other fields with less stringent criteria for the quality of an idea have no hope of using transformationally creative programs.

4.2.4. Literature

Like in most other domains, models of explorative creativity are used in literature. Yet, literature is one of the least successful fields of artificial creativity. In order to write a meaningful and interesting text, one needs a tremendous (and mostly implicit) knowledge about the world, the cultural context and the target audience. Human authors don’t even notice the huge knowledgebase that they automatically use in their writings.

The importance of the knowledgebase becomes only apparent, if it is lacking, as in the case of computer programs that try to produce texts.

Therefore, all programs to date are not able to produce a text that could trick naïve judges in believing it was written by a human.

Either these programs produce texts that lack coherence, are totally confusing and don't make any sense, as the program misses the world-knowledge to produce such a text - or the programs write texts that are meaningful but so constrained, that these texts are completely dull and uninteresting to a human reader. As these programs have – again – no inherent criteria about how to write an interesting and intriguing yet meaningful story, they can't do it.

One of these programs is “MINISTREL”, a program that tries to write thrillers about a murder. MINISTREL bases on a script-schemata-approach, assigning roles and agents to slots (the murderer, the victim, etc.), producing endless variations of the same basic plot. Besides problems with style (it simply can't build tension), it also can't introduce new characters or anything like that.

The attempt to apply computer models of creativity to literature could be considered an utter failure – due to the lack of world-knowledge and valid formal criteria of “good literature” – if there were no successful applications in poetry.

Surprisingly, poems – the very embodiment of creativity for the common man – are very easy to produce by computer models of creativity. Contemporary programs can produce poems that are either equal or surpassing the creativity (as judged by humans) of poems that are generated by even the finest human poets. This tells much about the creative effort that is involved in poetry: Not too much. As valid criteria for good poetry are totally lacking, no one notices that the computer doesn't have any. As no large knowledge base is needed, the computer programs can arbitrarily produce sentences – the weirder, the better. In the case of poems, creativity is in the mind (or the interpretation) of the reader anyway. Humans are used to encounter totally senseless poetry – the meaning is only added by the interpretation, never contained in the poem itself. Therefore, it is very ironic, but not really surprising, that poets can be so easily replaced by machines...

4.3. Evaluation of the artificial creativity approach

Artificial creativity can be seen as a special instance of AI research, an application of AI principles to creativity and a genuinely interdisciplinary approach between AI, psychology and the respective domain that is to be modeled (art, science, etc).

After our brief review of computer models of creativity – what can we learn about creativity from the artificial creativity approach and besides – how valid is the artificial creativity approach in the first place?

Interestingly enough, compared to artificial intelligence with all its problems, artificial creativity seemed to be relatively easy. Most of the philosophical ranting that plagues AI could be avoided for artificial creativity. Moreover, most programs of artificial creativity pass the “Turing-Test” for creative performances. Human judges can’t decide which creative performance is human and which is artificial.

This is pretty impressive and largely owed to a better conceptual basis than AI: While no one can agree on a formal definition of intelligence, the formal definition of creativity (novel and good) is widely accepted and serves as a starting point for all kinds of modeling approaches. In that sense, the conditions are far better, than in AI itself.

The remaining problems are largely owed to the fact that this formal definition of creativity has to be filled with flesh: “What kind of novelty do we want – combinatory, explorative or transformational?”; “How do we determine the value of an idea?”

These are hard problems and limit the practical use of computer models of creativity.

Nevertheless, the relative success of the different models in the different domains allows us to infer at least the following important aspects of the principles of (human) creativity in general, some of which are counter-intuitive:

- Obviously, random processes play a role in creativity. The notion of “accidental discovery”, the random linking of previously existing but independent ideas can be fruitful and creative, if the program (or organism) has criteria to select (and pursue

further) the best fruits of these random processes. No program could achieve any creative product without at least some sort of random sub-process.

- The notion that creativity is moderately domain-specific and that different domains require different kinds of creativity is strongly supported by the computer models of creativity. A creative program for devising mathematical theorems or suggesting experiments works totally different than their counterparts that produce art, while sharing the same formal definition of creativity for their outputs (new and good). Different domains also allow for different kinds of creativity, namely transformational vs. mere explorative.
- Strongly supporting the position of Weisberg, creativity is in fact bound to knowledge. Creative information processing requires much knowledge (as exemplified in literature). The better the knowledge-base (the richer the conceptual space), the better and more successful the creative outputs will be. Some domains require less knowledge – poetry, for example.
- Besides knowing the conceptual space, programs will be more creative, if they are persistent in the exploration of this space, supporting the notions of Simonton, who emphasizes productivity as a major contributor to creative achievements.
- The failure of many computer programs in many domains to achieve transformational creativity illustrates the fact, that rigid criteria for the evaluation of creative products are very important to allow for creativity. This fact has been overlooked in conventional research on creativity and defies public notions about the nature of creative thought. Moreover, the failure to provide formal criteria in most domains (except mathematics or engineering) supports the position of Csikszentmihaly who emphasizes the important role of the social context for creativity: He is right – the social context and the cultural value system provides most of the criteria that are

necessary for the evaluation of creative products. The criteria in most domains are not objective or formal, they are socially constructed.

In addition to these specific insights in the nature of creativity that are provided by computer models of creativity, there are some other general points that one can draw from this field:

- Artificial creativity can indeed be considered as an existence proof that creativity is in fact a computational problem and doesn't necessarily need to be explained in esoteric terms of "incubation", "illumination" or inspiration.
- This de-mystification of creativity might be a major bummer for all those romantics who like to think of creativity as a last resort for esoteric thought and free will, cherishing it as a paradise for human nature, an island within the ocean of heartless science. This state of affairs has ended with the advent of convincing computer models of creativity – thereby rightfully conquering this strange island for science. The fact that creativity is possible within fully deterministic and mechanical systems leads us to abandon the conventional idea of artificial and mechanistic/deterministic forces on one side and human and creative forces on the opposite. This postulated dimension doesn't exist, artificial creativity shatters these ill-framed but popular romantic notions.
- Computer models of creativity allowed further conceptual distinctions between different types of creativity, showing that exploratory creativity relatively easy and different from the paradigm-shifting transformational creativity.
- The major advantage of computer models of creativity remains the empirical testability of the models. The programmers have to be extremely explicit about every detail of their model – else, it won't work. This allows for a unprecedented conceptual sharpness in the field of creativity research, as the main drawback of most conventional theories of creativity is their notorious fuzzyness. This is not possible for computer models and this very fact is intriguing in itself.

- Besides informing us about the nature of creativity, future works in artificial creativity might focus on applications of these computer models, taking the same path, as AI. This would probably lead to programs that surpass human creativity in many if not all domains, allowing humanity to focus its manpower on other tasks.
- Speaking of applications, there are some tasks that artificially creative systems can already perfectly well do, allowing to release the respective humans into the labor force and the job-market. One prominent example are poets. Other artists like composers could follow soon thereafter. It makes perfect sense to replace these overrated professions of low-level creativity with machines.

Despite all these astounding advantages and successes of computer models of creativity, there remain some problems and open questions, most of which were already salient in previous sections. However, a quick summary of these issues is appropriate:

- As creativity is in fact fairly domain-specific, this might in principle require different „kinds“ of creativity for different domains - it's unclear, if an unified computational theory of creativity is even possible or not. This remains an open empirical question.
- At present, one of the most limiting factors for practical applications of computer models of creativity is the problem of how to enable the computer to evaluate its own novel "ideas" in a meaningful way. Until AI systems can be transformational creative in more domains in a fruitful way, their capacity to aid human thinking and problem solving will be strictly limited. The fact that AI and AC systems can't do it begs the question, why not. For once, social feedback that can serve as a value system is lacking. Another human value system – emotions – that play a major role in decision making and problem solving (see Damasio, 1995) is also missing. Maybe this problem can only be resolved if we can come up with the field of "artificial emotionality". From a scientific point of view, this would even be better, forcing the romantics out of another contemporary hideout. Maybe breakthroughs in machine

learning would also help to resolve this problem. How do humans evaluate things? They have value systems that derive from their socio-cultural environment and from phylogenetic (evolutionary) and ontogenetic learning, mediated by emotions. As computer models of creativity do none of these things, it is not surprising that they lack the ability to evaluate their products in a meaningful way.

- Certain – social – aspects of creativity might never be simulated by a computer, as these just are not mere intrapsychic processes. This is not a drawback, just something that has to be kept in mind to avoid conceptual dead-ends. This won't be possible and is not even necessary. After all, creativity is a highly complex construct with a very complex and sometimes (seemingly) contradictory phenomenology.

In conclusion, computer models show at the very least, there is – contrary to popular belief – probably nothing divine or indetermined about creativity. A computer is as fully rule-governed, mechanistic and determined as it gets, yet able to produce genuinely creative products. Of course, Turing's question, if a machine can be truly intelligent also applies to creativity: Can a machine be truly creative? Many people would still never go so far, as to grant a machine “intelligence” – with no real argument whatsoever. Now, we will probably have to accommodate the notion that machines could even be creative. The trend is still intact: The more science progresses, the less mysterious, esoteric and unique human behavior seems. Creativity can be conceptualized as a computational problem and it makes sense that humans also do it in a computational fashion. Yet, the computational problems are so complex that we invented funky terms like “inspiration” for these processes, just to cover our ignorance. The time has come to get rid of these preliminary and misleading concepts.

It won't hurt us to leave the question, whether a program or computer can ever be truly creative (as “truly intelligent”) for philosophers to ponder and will be essentially dependent on the respective definition. If the definition is merely formal, it will be possible to answer this question with a “yes”.

5 Outlook and conclusions

For the naïve and unarmed mind, creativity seems to be a deep mystery; that's, why people use all these weird and mysterious romantic concepts that seem satisfactory to the unarmed mind, yet without explaining anything. For the scientist, it's still hard to conceptualize, how something new and useful can be produced by the brain/mind in the first place.

Computer models of creativity help us to alleviate this condition: They exemplify, how a perfectly determined system can produce novel and valuable solutions to problems and how creative information processing COULD be done in principle. They also point out, how much skill and knowledge is in fact necessary to come up with truly creative work, illustrated by the still failing models in literature. Computer models of creativity also support the notion that there are no inherently "special" processes involved, as suggested by the psychoanalytic view on creativity. In addition, computer models of creativity powerfully demonstrate that creativity is not so mysterious after all and that it is a (even though extremely complex) computational problem that can be dealt with by algorithmic methods. Counter-intuitively, creativity is easier and more doable, if the problem is well-structured and the problem-space clearly constrained, both to explore and to transcend it. That the human mind and brain is capable of solving this kind of huge computational problem better than any computer models shows, how flexible, powerful and formidable the brain/mind-machine actually is.

In the long run, with the creation of a cognitive neuroscience of creativity, we will even be able to observe these computational principles in the generation of creativity in the brain itself. Biology will also place reasonable constraints on our computational models of creativity, forging a coherent science of creativity, including artificial creativity research. Reviewing the literature on creativity so far, one wonders if creativity is an issue for cognitive psychology in the first place – it certainly is an issue for psychology as a whole, but it's also certainly not ONLY an issue of cognitive psychology, more like an interdisciplinary approach within psychology, encompassing all subfields of psychology.

Besides and beyond that, very good models of creativity that possibly even surpass human creativity in some domains could assist and help human efforts. In science, for example, we could have the program read the literature and propose the experiments, while letting the undergrads conduct the study, so we could all just relax a little.

But in earnest: If we are able to model outstanding examples of human creativity in different fields like Einstein in science or Picasso in art, we will arrive at incredibly important insights about human nature in general and the nature of creativity in particular. This prospects should justify any effort to do research on questions like this, no matter how hard and complicated they might be. Clearly, creativity is an extremely hard problem – so we should finally take it seriously and not leave it to romantic dreamers, lunatics and scam-artists. The scientific neglect of creativity is currently still the case - as quick look in the local bookstore under the section “books on creativity” will easily reveal. Not many scientific journals are devoted to the study of creativity, yet thousands of bogus folk-psychological books and myths are easily available to the general public. This state of affairs is a pitiful shame that calls for radical change. Creativity is a fascinating phenomenon that can be studied on many levels of analysis, but – because of it’s complexity and many pitfalls – still a widely neglected field of research within science in general and psychology in particular.

Why should we leave the field to these people that parasite on the ignorance of the masses?

I can only renew the plea that Guilford already made in his presidential address over 50 years ago and call for more research on creativity. The tools to do that are about to be ready, waiting in computer science and neuroscience just to be used to resolve theoretical and empirical questions on creativity. Our basic understanding of creativity is still very poor. Almost nothing is resolved or presentable on a textbook-level. It is not even clear, if the persisting problems are conceptual or empirical problems. Again: We use the term creativity far more often than is justified by our understanding of this concept. Instead of leaving the field to speculation, this has to change. Let’s take this task seriously!

6 References

- AMABILE, T.M. (1985). Motivation and creativity: Effects of motivational conceptualization. *Journal of Personality and Social Psychology*, 48, 393-399.
- BODEN, M.A. (1992). *The creative mind: Myths and mechanisms*. London: Weidenfeld & Nicolson.
- BODEN, M.A. (1999). Computer Models of Creativity. In: Sternberg, R. J. (Ed.), *Handbook of Creativity* (pp. 351-372). Cambridge: Cambridge University Press.
- BODEN, M.A. (1990). *The Philosophy of Artificial Intelligence*. Oxford: Oxford University Press.
- BODEN, M.A. (1996). *Dimensions of Creativity*. Cambridge: MIT Press.
- BODEN, M.A. (2000). State of the art: Computer models of creativity. *Psychologist*, 13, 72-76.
- CAMPBELL, D.T. (1960). Blind variation and selective retention in creative thought as in other knowledge processes. *Psychological Review*, 67, 380-400.
- COHEN, H. (1981). *On the Modelling of Creative Behavior*. Santa Monica: Rand Corporation.
- CSIKSZENTMIHALY, M. (1996). *Creativity*. New York: HarperCollins.
- DAMASIO, A.R. (1995). *Descartes' Error : Emotion, Reason, and the Human Brain*. New York: Avon books.
- FELDMAN, D.H. (1988). *Creativity: Dreams, insights, and transformations*. In R. Sternberg (Ed.), *The nature of creativity* (pp. 271-297). Cambridge University Press.
- FINKE, R.A., WARD, T.B., & SMITH, S.M. (1992). *Creative cognition: Theory, research and applications*. Cambridge: MIT Press.
- GALOTTI, K.M. (1998). *Cognitive Psychology In and Out of the Laboratory*. Belmont: Wadsworth.
- GUILFORD, J.P. (1950). Creativity. *American Psychologist*, 5, 444-454.
- GUILFORD, J.P. (1967). *The nature of human intelligence*. New York: McGraw-Hill.
- LUCHINS, A.S. (1942). Mechanization in problem solving. *Psychological Monographs*, 54, 248.
- METCALFE, J., & WIEBE, D. (1987). Intuition in insight and non-insight problem solving. *Memory and Cognition*, 15, 238-246.
- NEWQUIST, H.P. (1994). *The brain makers*. Indianapolis: Sams.
- OSBORN, A.F. (1953). *Applied imagination*. New York: Scribner's.

- PERKINS, D.N. (1981). *The mind's best work*. Cambridge, MA: Harvard University Press.
- SIMONTON, D.K. (1988). *Scientific genius. A psychology of science*. New York: Cambridge University Press.
- SIMONTON, D.K. (1997). Creative productivity: A predictive and explanatory model of career trajectories and landmarks. *Psychological Review*, 104, 66-89.
- STERNBERG, R.J. (1988). *The triarchic mind: A new theory of human intelligence*. New York: Viking.
- STERNBERG, R.J. (1999). *Handbook of Creativity*. Cambridge: Cambridge University Press.
- SULLOWAY, F.J. (1997). *Born to rebel: Birth Order, Family Dynamics and Creative Lives*. New York: Vintage.
- THAGARD, P. (1999). *Kognitionswissenschaft – Ein Lehrbuch*. Stuttgart: Klett-Cotta.
- TORRANCE, E.P. (1974). *Torrance tests of creative thinking*. Bensenville, IL: Scholastic Testing Service.
- TURING, A.M. (1950). Computing machinery and intelligence. *Mind*, 59, 433-460.
- WALLAS, G. (1926). *The art of thought*. New York: Harcourt, Brace.
- WALLISCH, P. (2001). The cognitive neuroscience of attention – Happily ever after? *Unpublished term paper*.
- WEISBERG, R.W. (1986). *Creativity: Genius and other myths*. New York: W.H. Freeman.
- WEISBERG, R.W. (1993). *Creativity: Beyond the myth of genius*. New York: W.H. Freeman.
- WINSTON P.H. (1992). *Artificial intelligence*. Reading: Addison-Wesley.