



Towards creative information exploration based on Koestler's concept of bisociation

Dubitzky, W., Kötter, T., Schmidt, O., & Berthold, M. (2012). Towards creative information exploration based on Koestler's concept of bisociation. In *Bisociative Knowledge Discovery* (pp. 11-32). Springer.

[Link to publication record in Ulster University Research Portal](#)

Published in:

Bisociative Knowledge Discovery

Publication Status:

Published (in print/issue): 01/01/2012

Document Version

Publisher's PDF, also known as Version of record

General rights

Copyright for the publications made accessible via Ulster University's Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Ulster University's institutional repository that provides access to Ulster's research outputs. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact pure-support@ulster.ac.uk.

Towards Creative Information Exploration Based on Koestler's Concept of Bisociation

Werner Dubitzky¹, Tobias Kötter², Oliver Schmidt¹, and Michael R. Berthold²

¹ University of Ulster, Coleraine, Northern Ireland, UK
w.dubitzky@ulster.ac.uk, schmidt-01@email.ulster.ac.uk

² University of Konstanz, Konstanz, Germany
{tobias.koetter,michael.berthold}@uni-konstanz.de

Abstract. *Creative information exploration* refers to a novel framework for exploring large volumes of heterogeneous information. In particular, creative information exploration seeks to discover new, surprising and valuable relationships in data that would not be revealed by conventional information retrieval, data mining and data analysis technologies. While our approach is inspired by work in the field of computational creativity, we are particularly interested in a model of creativity proposed by Arthur Koestler in the 1960s. Koestler's model of creativity rests on the concept of *bisociation*. Bisociative thinking occurs when a problem, idea, event or situation is perceived simultaneously in two or more "matrices of thought" or domains. When two matrices of thought interact with each other, the result is either their *fusion* in a novel intellectual synthesis or their *confrontation* in a new aesthetic experience. This article discusses some of the foundational issues of computational creativity and bisociation in the context of creative information exploration.

"Creativity is the defeat of habit by originality." – Arthur Koestler

1 Introduction

According to Higgins, *creativity* is the process of generating something new that has value [19]. Along with other essentially human abilities, such as intelligence, creativity has long been viewed as one of the unassailable bastions of the human condition. Since the advent of the computer age this monopoly has been challenged. A new scientific discipline called *computational creativity* aims to model, simulate or replicate creativity with a computer [7]. This article explores the concept of *bisociation* [20] in the context of computational creativity. While our discussion may be relevant to a large number of domains in which creativity plays a central role, we emphasize domains with clear practical applications, such as science and engineering. We start our discourse on bisociation with the familiar concept of *association*.

The concept of association is at the heart of many of today's powerful computer technologies such as information retrieval and data mining. These technologies typically employ "association by similarity or co-occurrence" to locate or discover

information relevant to a user’s tasks. A typical feature of these approaches is that the underlying information pool (document corpora, databases, Web sites, etc.) contains information that has been pre-selected in some way to focus and simplify the discovery process. For example, a biological study would pre-select scientific papers from relevant life science journals or abstracts before applying a particular text mining task. Pre-selecting information in this way already introduces certain limits on how creative these conventional approaches can be. This means that under normal circumstances such resources would not be combined to facilitate creative insights and solutions. A novel information exploration paradigm that aims to facilitate the generation of *creative* insight or solutions could be referred to as *creative information exploration (CIE)*. Domains where CIE is critical include design and engineering, the arts (e.g., painting, sculpture, architecture, music and poetry) as well as scientific discovery disciplines.

In the remainder of this article we use the terms creative domains and creative disciplines to designate domains and disciplines in which creative information discovery plays an important role.

People working in creative domains employ creative thinking to connect seemingly unrelated information, for example, by using metaphors, analogy and other ways of thinking and reasoning [6]. Creative styles of thought allow the mixing of conceptual categories and contexts that are normally separated. Our goal is to develop computer-based solutions that support creative thinking. Inspired by Koestler’s notion of *bisociation* [20], our particular aim is to develop concepts and solutions that facilitate bisociative CIE tasks in creative domains. Intuitively, bisociative CIE could be viewed as an approach that seeks to combine elements from two or more “incompatible” concept or information spaces (domains) to generate creative solutions and insight.

The remainder of this article is organized as follows: Sections 2 and 3 introduce a working definition of creativity with a view to its computational realization. In Sections 4 to 6 we review Koestler’s notion of bisociation and offer an initial formal definition of this concept. Before we reflect on the work presented in this article and offer some concluding remarks (Section 8), we present a short review of related work in Section 7.

2 Creativity

2.1 What Is Creativity?

Human creativity, like other human capabilities, is difficult to define and formalize. In this article we adopt the following working definition of *creativity* based on the work by Margaret Boden [6].

Definition 1 (creativity). *Creativity is the ability to come up with ideas or artifacts that are new, surprising, and valuable.*

In this working definition of creativity the notions of *idea* and *artifact* refer to concepts and creations from art as well as science and engineering and other areas. Here we view creativity as an ability which is an intrinsic part of an

intelligent agent (human, machine-based or otherwise). In the following discussion we elaborate the meaning of the concepts *new*, *surprising* and *valuable* in the definition of creativity.

The word *new* in our working definition of creativity may refer to two dimensions: *historic creativity* or *personal creativity*. By historic creativity we mean ideas or artifacts that are original in the sense that they represent the first occurrence of a particular idea or artefact in human history. The history of science and modern legal practice tell us that sometimes it may not be straightforward to determine precisely the first occurrence of a scientific or engineering idea. Examples of disputes over historic creativity include the theory of evolution, the invention of gun powder, and the social Web site Facebook. Personal creativity, on the other hand, means that someone comes up with an idea or invention independently from someone else who had already conceived of the same thing *before*. From the perspective of the “re-inventor” this still constitutes “true” creativity.

An important factor in our working definition of creativity concerns the notion of *surprise* – for a new idea to be considered creative there has to be an element of surprise. An idea or artefact may be surprising because it is unlikely (has a low probability of occurring) or unfamiliar. When a new idea unexpectedly falls into an already familiar conceptual framework (or thinking style) one is intrigued to not have realized it before. For example, in 1996 Akihiro Yokoi invented a “digital pet” called Tamagotchi which soon became a best seller. While the concept of looking after plants, pet animals and soft toy pets has been around for a long time, no one had dared to think that this idea could be applied to devices that resemble digital pocket calculators. A different type of surprise occurs when we encounter an apparently *impossible* concept or artefact. For instance, in 1905 Einstein shocked the scientific establishment by suggesting that energy is being *transmitted* in finite “packets” called quanta [11]. Max Planck, the originator of quantum theory, initially rejected Einstein’s proposal even though his own theory suggested that energy *transfer* to and from matter is not continuous but discrete.

The last element in our working definition of creativity is the notion of *value* – a new concept or artefact must be valuable in some non-trivial way to qualify as creative. In the fine arts aesthetic values are difficult to recognize or agree about: what makes a painting by one artist hundred times more expensive than a painting by another? To formally define aesthetic values is even harder. Furthermore, values vary over time and within and across cultures. Even in science there is often considerable disagreement over the “simplicity”, “elegance” or “beauty” of a theory or scientific argument. Einstein and Bohr, for instance, had argued over decades about the value (correctness and completeness) of the two prevailing models of the atom (the probabilistic and discrete model, favored by Bohr, and the deterministic and continuous model, which was preferred by Einstein) [22]. Whether a particular hypothesis is interesting or valuable may depend on scientific, social, economic, political and other factors. So even when we agree on novelty and the factor of surprise, there may still be a considerable disagreement over how valuable a new idea or artefact is, hence over the degree of creativity.

This brief discussion about the nature of creativity and the difficulty to recognize and agree on what creativity actually is serves as a context for the development of computational creativity techniques. Ultimately, what constitutes human or machine creativity is difficult to judge and needs to be assessed on a case-by-case basis.

2.2 Three Roads to Creativity

Following Boden [6] we distinguish three processes of creativity; these relate to the three forms of surprises discussed above.

Combinatorial Creativity. Arthur Koestler¹ is credited with the following characterization of creativity:

The creative act is not an act of creation in the sense of the Old Testament. It does not create something out of nothing; it uncovers, selects, reshuffles, combines, synthesizes already existing facts, ideas, faculties, skills. The more familiar the parts, the more striking the new whole.

This idea is very much in line with the first process of creativity identified by Boden, which generates unfamiliar combinations of familiar concepts and constructs. In humans, analogy is a fundamental cognitive process in which familiar elements appear in an unfamiliar arrangement. A typical example of analogy establishes an analogical relationship between Niels Bohr’s model of the atom with the basic structure of the heliocentric solar system. Facilitating this kind of creative process requires a rich knowledge structure and flexible ways of manipulating this structure. Clearly, the novel combination of elements must have a point or a meaning. Therefore, purely random shuffling and re-combination of elements will not be sufficient to generate creativity.

Exploratory Creativity. Margaret Boden defines *conceptual spaces* as a “structured style of thought”. In her definition, a key characteristic of conceptual spaces is that they are not originated by an individual but are a structure adopted from the cultures and peer groups within which people live [6]. Conceptual spaces include ways of writing prose, styles of architecture and art, theories of nature, as well as approaches to design and engineering. So any systematic way of thinking which is valued by a certain group or culture could be thought of as a conceptual space.

In Boden’s framework, a conceptual space defines a space of possible combinations of its elements, where each combination represents a particular thought, idea or artifact. While the number of possible thoughts within a conceptual space may be very large, only a fraction of these may have actually been realized. Consider, for instance, the games of chess and checkers. In chess the number of possible legal positions or “configurations” has been estimated at $10^{15\ 790}$ and for checkers the number is 10^{18} [16,29]. Clearly, even with the long history of chess playing, only a very small number of possible “combinations” could have

¹ Prolific writer and author of *The act of creation* [20].

been explored so far. Clearly, Boden's concept of a conceptual space is much broader. For the game of chess, for example, it would not only include all possible chess board positions but also all knowledge structures employed by chess players to play the game as well as other facts and information about chess.

No matter what the actual size of a given conceptual space, someone who comes up with a new combination within that space is considered to be creative in the exploratory sense (provided the combination "has a point"). Boden likens the exploration of conceptual spaces to the exploration of a territory with a map. The map encompasses all possibilities, but to discover a particular and valuable possibility one needs to go out and explore the actual territory. *Exploratory creativity* is important as it facilitates the discovery of so far unknown possibilities. Once such novel possibilities come to light, the explorers may even be able to reflect deeper on the limits and potentials of a particular conceptual space.

Transformational Creativity. Exploratory creativity is limited by the possibilities defined within a conceptual space or thinking style (or "map"). Essentially, each conceptual space restricts the kind of thoughts that can be thought. To overcome this limitation, and to attempt to think what is unthinkable within a given conceptual space, it is necessary to change or *transform* the conceptual space. It must be transformed so that thoughts that were inconceivable within the previous version of the space now become possible. Such transformations may be subtle or radical. *Transformational creativity* constitutes the deepest form of creative processes in Boden's model of creativity.

3 Computational Creativity

Teaching humans to be creative is a flourishing business and the number of creativity techniques available is large [19]. Teaching or programming a computer to be creative or appear to be creative is another matter altogether. *Computational creativity* refers to an active scientific discipline that aims to model, simulate or replicate creativity using a computer [7].

Computational creativity draws on many concepts developed within the field of *artificial intelligence* (AI). Analogously to computational creativity, AI could be defined as a discipline aiming to model, simulate or replicate (human) *intelligence*. Boden suggests that AI concepts could be used to define and construct artificial conceptual spaces which could then be studied and eventually be used to combine elements from the spaces, and to explore and transform such spaces with the aim of generating creative insight and solutions. Boden describes concrete AI-based approaches to computational creativity [6,7].

4 Koestler's Concept of Bisociation

People working in creative domains employ creative thinking to connect seemingly unrelated information (true negatives under the association paradigm), for example, by using a metaphoric or analogical way of thinking. Analogical and metaphoric styles of thought allow the mixing of conceptual categories and

contexts that are normally separated. In the 1960s Arthur Koestler developed a model of creative thinking referred to as *bisociation* [20]. Bisociation facilitates the mixture in one human mind of concepts from two contexts or categories of objects that are normally considered separate by the literal processes of the mind.

Koestler proposed the bisociation concept to distinguish the type of metaphoric thinking that leads to the acts of great creativity from the more “pedestrian” associative style of thinking, with which we are so familiar in our everyday lives and which pervades many of today’s computing approaches. Associative thinking is based on the “habits” or set of routines that have been established over a period of time. Associative processes combine elements from the same “matrix” of thought. The associative mode of thinking differs from the bisociative mode that underlies the creative act. Bisociation, according to Koestler, means to join unrelated, often conflicting, information in a new way. It is being “double minded” or able to think simultaneously on more than one plane or matrix of thought (see Figure 1). “When two independent matrices of perception or reasoning interact with each other the result ... is a ... fusion in a new intellectual synthesis ...” [20]. Frank Barron reinforces this idea and characterizes bisociation as “the ability to tolerate chaos or seemingly opposite information” [3]. Koestler makes a clear distinction between more routine or habitual thinking (association) operating within a single plane or matrix of thought, and the more creative bisociative mode of thinking which connects independent autonomous matrices.

Koestler’s basic concept of bisociation is illustrated in Figure 1. The diagram depicts two matrices of thought (domains or knowledge bases in our terminology), M_1 and M_2 , as orthogonal planes. M_1 and M_2 represent two self-contained but “habitually incompatible” matrices of thought. An event, idea, situation, concept or problem, π , which is perceived simultaneously in both matrices is not merely linked to one associative context (M_1 or M_2) but *bisociated* with two associative contexts (M_1 and M_2). In the diagram, π is illustrated by the thick line cutting across M_1 and M_2 . The diagram illustrates six concepts labeled c_1, \dots, c_6 . The concepts c_1, c_2, c_3 and c_6 are perceivable in matrix M_2 and c_1, c_2, c_3, c_4 and c_5 are perceivable in M_1 . The concepts c_1, c_2, c_3 are associated with the problem π – because c_1, c_2, c_3 are perceivable in both matrices, it is possible to “see” the problem simultaneously from two frames of mind.

Central to Koestler’s concept of bisociation are the notions of a *matrix* and a *code* Koestler [20]; we quote from page 38:

... to introduce a pair of related concepts which play a central role in this book and are indispensable to all that follows. ... I shall use the word ‘*matrix*’ to denote any ability, habit, or skill, any pattern of ordered behavior governed by a ‘*code*’ of fixed rules.

A matrix² in Koestler’s framework denotes any ability, skill, habit or pattern of ordered behavior. Matrices shape our perceptions, thoughts, and activities; they

² Other terms Koestler uses for the concept of a matrix include the following: matrix of thought, matrix of behavior, matrix of experience, matrix of perception, associative context, frame of reference, universe of discourse, type of logic, code of behavior.




Fig. 1. Illustration of Koestler's concept of bisociation (adapted from Koestler [20])

could be viewed as condensations of learning into “habit”. For example, a spider has an innate skill that enables it to build webs, a mathematician possesses the ability of mathematical reasoning, and a chess grandmaster has a knowledge base which allows him to play chess at a very high level. The abilities and skills represented by a matrix may be applied to concrete problems and tasks in a flexible way. For example, depending on the environment a spider finds itself in, it may choose three, four or more points of attachment to suspend its web.

Each matrix in Koestler's model of bisociation is governed by a set of fixed *codes* or rules. The rules could be innate or acquired. For example, in the game of chess, the rules of the game are fixed, while the patterns of knowledge (allowing one to play well or not so well) vary across players³. In mathematics, operations such as multiplication, differentiation, integration, etc. constitute fixed rules that govern mathematical reasoning. Another example of a code are the assumptions, concepts, notions, etc. that underly religious, political, economic, philosophical and similar debates and arguments. For instance, a debate on abortion may be held “in terms of” religious morality or social responsibility. Often the rules that govern a matrix of skill (ability, habit) function on a lower level of awareness than the actual performance itself (playing the piano, carrying out a conversation, formulating a strategy).

Once people have reached adulthood they have formed more or less rigid, automated patterns of behavior and thinking (“habits” or knowledge bases). Sometimes these patterns are interrupted by spontaneous sparks of insight which presents a familiar concept or situation in a new light. This happens when we connect previously unconnected matrices of perception or experience in a creative act of bisociation. Considering the field of humor, science and engineering as well as the arts, Koestler's conjecture was that bisociation is a general mechanism

³ Certain ways of playing chess are also relatively frequent or nearly constant. For example, certain moves in chess openings, or certain endgame patterns.

for the creative act. When two habitually independent matrices of perception or reasoning interact with each other the result is either a *collision* ending in laughter, or their *fusion* in a new intellectual synthesis, or their *confrontation* in an aesthetic experience [20].


Koestler provides numerous examples and illustrations of his bisociation concept in different areas and domains. In the following we briefly summarize the Archimedes example which Koestler refers to as the “Eureka act” (Figure 2). The Eureka act is concerned with the discovery of solutions to a more or less scientific problem.

Archimedes, a leading scientists in classical antiquity, was tasked with the problem of determining whether a crown (a present for Hiero, tyrant of Syracuse) consisted of pure gold or was adulterated with silver. To solve this problem Archimedes needed to measure the volume of the crown. At the time no method existed to determine the volume of such an irregularly shaped three-dimensional object. Pondering over this problem, Archimedes’s thoughts wandered around his matrix of geometrical knowledge (Figure 2a). One day, while taking a bath, Archimedes noticed the rise of the water level as his body slid into the basin. It was at this point when he connected the matrix of and experience associated with taking a bath with the matrix of his knowledge of geometry. He realized that the volume of water displaced was equal to the volume of the immersed parts of his own body. This Eureka moment is illustrated in Figure 2b. When Archimedes found the solution to this problem both matrices (associations of taking a bath and knowledge of geometry) were simultaneously active. In a sense Archimedes was looking at the same problem from two different perspectives of knowledge or experience at the same time. This “double-mindedness” allowed him to see the solution which was obscured under the view of either of the two individual perspectives.

Consider the diagram in Figure 2a. The dashed line illustrates Archimedes’s search through the conceptual space to find a solution for his problem. While the search path traverses both knowledge bases (M_1 and M_2), the reasoning of Archimedes is initially confined to perceiving only one knowledge base at a time. Thinking about the problem in this “habitual” way, Archimedes fails to “see” the solution, because he does not simultaneously perceive the concepts describing the solution (c_1 and c_6) and the problem (c_1 , c_2 and c_3).

Now consider the diagram in Figure 2b. At some point Archimedes is able to perceive the concepts describing both the problem (P) and the solution (S) simultaneously from the perspective of both knowledge bases. This is depicted by the line connecting the corresponding concepts across both knowledge bases. It is at this point when Archimedes experiences the Eureka moment which is created by the bisociative view spanning two matrices of thought.

The example of Archimedes and the crown illustrates how a familiar but unnoticed aspect of a phenomenon (rise of water level as a result of the immersion of an object) is suddenly perceived at an unfamiliar and significant angle (determining the purity of substance of an irregularly shaped object). Koestler refers to this as the “bisociative shock” often associated with discoveries when we suddenly see familiar objects and events in a strangely new and revealing light.



Legend:

- Concept as perceived within M_2
- Concept as perceived within M_1

Fig. 2. Illustration of the Eureka act (adapted from Koestler [20]). The matrix or knowledge base M_1 represents concepts or associations of *geometrical knowledge*, and M_2 those of *taking a bath*. The dashed lines represent the search or exploration of the matrices as part of the problem-solving process. (a) Diagram on the left: The line connecting the concepts c_1 , c_2 and c_3 represents the problem, P , as perceived by Archimedes based on his geometric knowledge base M_1 . The arc connecting the concepts c_1 and c_6 in M_2 represents the solution, S . (b) Diagram on the right: The concepts associated with the problem *and* solution when perceived *simultaneously* in both knowledge bases.

The distinguishing characteristics of associative and bisociative thought are summarized in Table 1.

Table 1. Comparison of characteristics of bisociation and association based on Koestler [20]

| Habit (Associative) | Originality (Bisociative) |
|---|-------------------------------------|
| association within a given matrix | bisociation of independent matrices |
| rigid to flexible variations on a theme | super-flexibility |
| repetitiveness | novelty |
| conservative | destructive-constructive |

5 Elements of Bisociative Computational Creativity

Before we formally define bisociation, we analyze and compare the concepts and models of creativity proposed by Boden and Koestler. We do this by adopting an AI perspective of the notions involved and, on this basis, attempt a synthesis. In essence, we define creativity and bisociation in terms of *domain theories* and *knowledge bases*. Simply put, a domain theory consists of all knowledge (concepts) relevant to a given domain (at a given point in time), regardless of the type of knowledge, how it is encoded (formalism, substrate) or where it is

located. Under this definition of a domain theory, a knowledge base is simply a subset of all the concepts in a domain theory. However, different classes of knowledge bases may be distinguished.

Concept. A *concept* denotes a cognitive unit of meaning which is sometimes referred to as “unit of knowledge”. Concept descriptions are constructed from concept *properties* (features, dimensions) [30]. A concept is normally associated with a corresponding representation or encoding in a particular language or formalism. Concepts form the basis for the cognitive abilities of an intelligent agent. Without a concept, an intelligent agent or reasoner, relying on a memory containing a potentially large number of items, would be hopelessly lost. If a reasoner perceived each entity as unique, it would be overwhelmed by the enormous diversity of what it experiences, unable to remember but a fraction of what it encounters. A concept captures the notion that many objects, ideas or events are alike in some important respects, and can therefore be thought about in the same or similar ways. Once an entity has been assigned to a concept on the basis of its perceptible properties, a concept may also be used to infer some of the entity’s non-perceptible attributes. Having, for example, chosen perceptible attributes like size, shape and material to decide an object is a book, it can be inferred that the object contains pages and textual information. This idea of inferability is based on the assumption that all instances of a given concept are subject to the same, or similar, underlying mechanisms (e.g., cause-effect relationships) which may or may not be completely known. Such mechanisms may be simple, in the case of books, or complex in chess positions.

Different views and models of concepts have been proposed [30]; these vary in a number of aspects, in particular, in the degree to which they are deterministic/probabilistic and intensional/extensional. In this article, concepts form the basic units from which domain theories and knowledge bases are constructed. Here, concepts include all forms of knowledge, including the three kinds of knowledge normally distinguished in epistemology: “knowledge that” (propositional, declarative knowledge), “knowledge how” (procedural knowledge) and “acquaintance knowledge” (about places, situations, cases, experiences) [1]. The knowledge that concepts represent may be tacit or explicit, it may be implemented on living tissue, electronic structures, paper or any other substrate. Critical for our discussion on domain theories and knowledge bases is that concepts are normally associated with one or more domains.

Notice that here we do not differentiate the representation languages or formalisms used to specify concrete knowledge structures (frames, rules, trees, networks, heuristics, case bases, etc.).

Domain Theory. For the purpose of this discussion, a *domain* is viewed as a formal or common sense topic, subject area, field of interest, for example, a scientific discipline (e.g., biology), a game (e.g., chess), social, cultural, economic or political topics (e.g., religious morality), common patterns of activity (e.g., taking a bath), and so on. Based on this view of a domain, we define a domain theory as follows:

Definition 2 (domain theory). *A domain theory D_i defines a set of concepts (knowledge units) that are associated with a particular domain i .*

Notice that a particular concept may belong to more than one domain theory at the same time.

In this view of a domain theory it is easy to see that most domain theories would be formed from an heterogeneous and distributed pool of knowledge “sources”, including humans, documents, electronic systems, and so on. For example, the domain theory of chess would be “encoded” in books, reports of tournaments, databases, chess programs, and the minds of a large number of chess players. While many of the concepts within the domain theory of chess would be shared across many chess players, other concepts may be unique to and accessible by individual players only (or by groups of players)⁴.

A domain theory is shared across a peer group. One consequence of the distributed and heterogeneous nature of most non-trivial domain theories is that they are usually associated with a particular peer group, culture, society, etc., rather than with an individual or a very small group of people. Notice, a domain theory, as it is defined here, usually includes elements that are not accessible by the entire peer group associated with it. For example, the subjective case base (acquaintance-knowledge) a particular chess master has accumulated over his career is not likely to be accessible by other chess masters (members of the peer group). Likewise, certain documents or electronic resources about chess knowledge may be accessible only to a limited group of peer members.

A domain theory is fixed or changing only very slowly. An established domain theory would normally not change radically but remain relatively stable and undergo mostly minor modifications over time. Radical changes of a domain theory would be related to changes in fundamental concepts of a domain theory. For example, nowadays in chess it rarely happens that a “standard” move in a particular game would be shown to be unsound.

A domain theory incorporates “hidden” concepts. At a given point in time, a large database holds facts and patterns that have already been explicitly reported or are known by at least one intelligent agent. However, at the same time there may be many “hidden” facts or patterns contained in the same database which have not been discovered yet. Analogously, a domain theory captures concepts which are explicitly documented or known by an intelligent agent. At the same time, a domain theory harbors concepts which are yet to be discovered. Notice that while the total number of hidden and explicit/known concepts within a domain theory may be very large (or even infinite), not every conceivable concept may be expressible under the constraints of a particular domain theory.

Knowledge Base. *A knowledge base is constructed from the concepts of a domain theory; we define a knowledge base as a subset of a domain theory as follows:*

⁴ Detailed psychological studies suggest that, for example, the number of symptom-illness correspondences known by a medical specialist, or the number of board positions memorized by a chess master, appear to be in the range of 30 000 to 100 000 [23].

Definition 3 (knowledge base). *A knowledge base K_i is defined as a subset of a domain theory D_i , i.e., $K_i \subseteq D_i$.*

This means that in the extreme case a knowledge base and a domain theory could be identical. This is of course only a theoretical possibility, because for real-world domain theories, a knowledge base is normally a highly selective subset of the domain theory. In particular, a knowledge base would tend to have the following characteristics.

A knowledge base is domain-specific. As a consequence of how a knowledge base is defined, it is always defined with respect to a particular domain. Hence, a knowledge base contains only concepts from the underlying domain theory.

A knowledge base is focused, selective, goal-oriented, biased ... A knowledge base is normally not formed by a random process which selects elements from a domain theory and puts them together to make up a knowledge base. Instead, a knowledge base is either intentionally constructed or it is evolved, and as a consequence a knowledge base normally represents a focused, selective, goal-oriented, biased, subjective, etc. subset of the domain theory. When a knowledge base is designed, its construction is guided by the function it is supposed to fulfill, by other design constraints and requirements, and by the set of biases, skills, abilities, etc. of its designers. In this process particular choices are made in terms of which concepts from the underlying domain theory will be included in the knowledge base. When an intelligent agent acquires knowledge (learning, evolution) it normally does so under a set of constraints, including the goals it pursues, its prior experience, abilities, skills, the environment it operates in, and so on. A knowledge base which is thus constructed or evolved has selected (or acquired) a set of domain concepts in a very biased or “habitual” way. Notice, as an intelligent agent evolves a knowledge base, it does not only assimilate knowledge from the domain theory that is shared by other peer members, but it also creates a part of the domain theory space that is normally not accessible to other peer members of the domain.

Agent-specific knowledge bases. In our definition of a knowledge base, a book on a particular variant of the Sicilian Defence could be considered as a knowledge base in the domain of chess. Often, however, in this discussion we are concerned with knowledge bases that are tied to or integrated within a specific intelligent agent⁵. In this case, we are talking about the type of knowledge base which is highly subjective, containing domain concepts which are not shared with the domain’s peer members. It is precisely the non-shared concepts in such an agent-specific knowledge base that form a kind of “inertial system” or “reference system” against which the common or shared parts of the knowledge base are viewed and interpreted. What is important to understand is that an intelligent agent has *exactly one* knowledge base for a given domain! This knowledge base may be

⁵ Here we use the term “intelligent agent” to denote a uniquely identifiable entity with cognitive abilities such as reasoning, planning, hypothesizing, etc. It is irrelevant on which physical substrate such an entity is implemented or whether or not it is highly localized in physical space.

empty, if the agent knows no concept in that domain, or it may be a non-empty knowledge base consisting of shared and non-shared concepts of the domain theory. The non-shared domain concepts impose a unique, biased perspective of the agent on the domain. The fact that an agent captures part of the domain theory which is normally not shared with other agents in the domain, makes such an agent-specific knowledge base special.

Agent-specific knowledge bases are habitually incompatible. Another critical aspect of the concept of an agent-specific knowledge base is that, given a concrete problem, normally (or “habitually”) only a single knowledge base would be active at a given time. This is what Koestler refers to as “habitually incompatible” matrices.

Models of Creativity. Both Boden and Koestler base their models on a corpus of domain-specific knowledge or concepts called *conceptual space* by Boden and *code* by Koestler. In our conceptualization both a code and a conceptual space are viewed as a *domain theory*.

With respects to Boden's model of creativity, domain theories are equivalent to the notion of *conceptual spaces*. They satisfy the characteristic of not being tied to an individual as well as being relatively stable over time. Indeed, a domain theory encompasses all the knowledge (or concepts) known about a domain at a given point in time. Furthermore, a domain theory represents Boden's “generative structure” [6] that contains the “possibilities” of hitherto unknown knowledge which may be discovered in the creative process (combinatorial or exploratory creativity). Essentially, these are all possible concepts within a domain theory that have not been made explicit in any form (documented) or are not known by any agent of the domain's peer group. Boden's transformational creativity is facilitated by a change or transformation of the underlying domain theory. Such a change would typically be realized by a modification or addition of concepts in a given domain theory.

In Koestler's framework of creativity the notion of a *code* is equivalent to our concept of a domain theory. Like Koestler's concept of a code, a domain theory constitutes a relatively fixed system of rules (or concepts) which governs the processes of creativity.

Unlike Koestler's model, which incorporates the notion of a *matrix*, Boden does not make a distinction between a matrix and a conceptual space. Comparing her model with that of Koestler, Boden states: “Matrices appear in my terminology as conceptual spaces, and different forms of bisociation as association, analogy, exploration, or transformation.” [6]. This is where Koestler's model appears to be more differentiated. With the notion of a matrix, Koestler puts the subjective perspective of the entity that engages in creative thought in the center of his model. Indeed, the matrix notion provides this degree of individuality that appears to be associated with many creative ideas and inventions. In our model, Koestler's matrix concept is reflected in the concept of a knowledge base. A knowledge base, like a matrix in Kostler's framework, is uniquely linked to a particular reasoner or intelligent agent. Indeed, a knowledge base carries the characteristics that Koestler associates with his matrices:

1. There is exactly one knowledge base per agent for each domain.
2. A knowledge base reflects the subjective personal, prejudiced and unobjective views and patterns of thinking and behavior – i.e., a *habitual frame of thought* – that provide a unique (albeit biased) perspective of the domain. Usually, when pondering over a task or problem, only the concepts of a single knowledge base would be active. This is why Koestler calls his matrices “habitually” incompatible. This notion does not seem to be reflected in Boden’s model.
3. Because each agent or reasoner incorporates a set of (partially overlapping) knowledge bases in a *highly integrated* fashion (with in a single “mind”), such an agent is equipped with the unequaled potential to discern patterns of bisociation by bringing together or superimposing multiple knowledge bases simultaneously. It is this structure that allows an agent to “see” or perceive a problem, situation or idea simultaneously from different frames of mind (knowledge bases).

Viewing Koestler’s matrix as a knowledge base appears to be a more realistic model for combinatorial, exploratory and transformational creativity, because it takes into account the fact that an entity’s (agent) view of the world is normally limited by the set of knowledge bases it has. One can assume that agents operating on the basis of Boden’s conceptual spaces are also limited to a subset of the conceptual space, but this is not so clear in the model of Boden.

Boden argues that bisociation can be incorporated in her model. However, in the absence of a clear account of the “habitual” dimension (represented by matrices in Koestler’s framework and by knowledge bases in our model) involved in bisociation, Boden’s model seems less convincing.

6 Towards a Formal Definition of Bisociation

Based on above considerations we now attempt to provide a formal definition of bisociation. In our definition we employ the following symbols:

Let U denote the *universe of discourse*, which consists of all concepts.

Let $c \in U$ denote a *concept* in U .

Within the universe of discourse, a problem, idea, situation or event π is associated with the concepts $X \subset U$. Typically, in a concrete setting, a subset $P \subset X$ is used to describe and reason about π .

D_i denotes a *domain theory* which represents the total knowledge (concepts) within a domain. Notice that the union of all domain theories represents the universe of discourse: $\cup_i D_i = U$. Furthermore, $\exists i, j : D_i \cap D_j \neq \emptyset$. This means that many domain theories overlap.

R denotes a *reference system* or *intelligent agent* which possesses exactly one knowledge base (empty or non-empty) per domain theory D_i .

$K_i^R \subset D_i$ denotes the *knowledge base* with respect to the reference system or intelligent agent R and domain theory D_i . Notice, an intelligent agent R has exactly a single knowledge base K_i^R (empty or non-empty) per domain theory i . For example, the knowledge base K_{chess}^R defines the chess knowledge base an intelligent R has.

$K^R = \cup_i K_i^R$ denotes the entire *set of knowledge bases* incorporated in the reference system or intelligent agent R . K^R represents the total knowledge that R has in all the domains. For example, an intelligent agent R may possess non-empty knowledge bases for the domains of chess, biology and religious morality, and an empty knowledge base for the domain of geometry.

Definition 4 (habitually incompatible knowledge bases). *Two agent-specific knowledge bases K_i^R and K_j^R ($i \neq j$) are said to be habitually incompatible if, at a given point in time t , there is no concept $c : c \in K_i^R \wedge c \in K_j^R$ that is active or perceived simultaneously in K_i^R and K_j^R .*

In other words, an intelligent agent usually employs a single frame of mind (knowledge base) at a given moment in time to think about a problem. One could compare this “pedestrian” way of thinking to a “sequential” mode of reasoning in which a reasoner switches between the matrices (knowledge bases) but only uses one matrix at the time.

Definition 5 (bisociation). *Let π denote a concrete problem, situation or event and let $X \subset U$ denote the concepts associated with π . Further, let K_i^R and K_j^R denote two habitually incompatible agent-specific knowledge bases ($i \neq j$). Bisociation occurs when elements of X are active or perceived simultaneously in both K_i^R and K_j^R at a given point in time t .*

This refers to the situation where a problem is perceived simultaneously in two frames of reference or matrices of thought (Figure 1).

For example, at time t the concepts $B = \{c_1, c_2, c_3\}$ may be active or perceived simultaneously in K_i^R and K_j^R . In this case we say that the concepts in A are *bisociated*.

Definition 6 (association). *Let π denote a concrete problem, situation or event and let $X \subset U$ denote the concepts associated with π . Further, let K_i^R denote an agent-specific knowledge base. Association occurs when elements of X are active or perceived in K_i^R at time t only.*

For example, at time t the concepts $A = \{c_1, c_2, c_3\}$ may be active in K_i^R only. In this case we say that the concepts in A are *associated* (with each other).

7 Related Work

The key notion of bisociation is a knowledge structure that is defined on the concepts originating from multiple domains. Below we briefly look at some of the literature which is closely related to bisociation. This short review does not claim to be exhaustive. A more comprehensive literature review should include areas such as data and information fusion, heterogenous information networks, interchange of knowledge bases and ontologies, multi-agent systems, hybrid intelligent systems, metaphor-based reasoning (conceptual/cognitive metaphors), conceptual blending, discourse reasoning, and others.

Analogical Reasoning. Analogy is a powerful form of logical inference which allows to make assertions about an entity or concept, X , based on its similarity with another entity or concept, Y . For example, we use our knowledge about water flow to determine properties of electrical circuits. The underlying assumption of analogical reasoning is that if two entities or concepts are similar in some respects, then they are probably alike in other respects as well. Like inductive reasoning, which proceeds from the particular to the general, analogical reasoning does not guarantee the truth of the conclusion given a true premise. Despite this similarity with inductive reasoning, analogical reasoning is often viewed as a form of reasoning which is distinct from inductive reasoning. For instance, Sowa and Majumdar view analogical reasoning as a two-step reasoning process which first inductively creates a theory from a set of cases, and then deductively generates an answer to a specific question or problem on the basis of the theory [32]. In AI, analogical reasoning is often described as a representational or *analogical mapping* from a known “source” domain to a (novel) “target” domain [17].

A key element in analogy is the mechanism of selection. Not all commonalities between two concepts are equally important when we compare the concepts and make predictions based on similarities. Therefore, a central issue in analogical mapping is to determine the selection constraints that guide our assessment of similarity and dissimilarity⁶. Two broad classes of selection constraints have been investigated in AI: goal-relevance and structure-relevance. The former is used to focus analogical mapping on information that is considered critical to the problem or goal at hand. The latter is used to guide analogical mapping based on the structural commonalities between two entities or concepts.

⁶ Similarity should consider the common and distinctive features of the entities under investigation. For example, let x and y denote two entities, and X and Y the sets of their characterizing features. Then the similarity, $sim(x, y)$, between x and y is a function of their common and distinctive features as follows:

$$sim(x, y) = \theta f(X \cap Y) - \alpha g(X \setminus Y) - \beta h(Y \setminus X),$$

where $f(X \cap Y)$ expresses the *similarity* based on common features in x and y , $g(X \setminus Y)$ the *dissimilarity* based on properties x has but y does not, and $h(Y \setminus X)$ the *dissimilarity* based on properties y has but x does not. θ , α and β influence how the various components affect the overall score, with $\theta, \alpha, \beta \in [0, 1]$.

Investigating the mechanisms of analogical reasoning in humans, Gentner and co-workers developed the *structure-mapping theory* of analogy [13]. The underlying assumptions in the structure-mapping theory are that (a) connected knowledge (concepts) is preferred over independent facts; this assumption is known as the systematicity principle, and (b) analogical mappings are based on structure-relevance selection constraints. The structure-mapping theory has been used to create a computational model called the *structure-mapping engine* [12]. The structure-mapping engine can find a mapping between the appropriate relations (between concepts in the considered domains) given a properly constructed representation of the domains of interest. Chalmers and co-workers [9] proposed a different approach to explain and model analogical reasoning. They view analogical reasoning as a product of a more general cognitive function called *high-level perception*. Morrison et al. interpret high-level perception and the structure-mapping theory as two aspects of analogy, rather than viewing them as mechanisms on two distinct cognitive organizational levels [27].

Human cognition is continually establishing *potential* mappings between knowledge domains or contexts. Analogical mapping occurs in a richly interconnected conceptual space in long-term memory. Attribute/category information plays a crucial role for the discovery of analogies across the conceptual spaces in long-term memory. Based on such a model of human memory, the following (simplified) analogical reasoning processes could be distinguished [14]:

1. **Retrieval:** In response to some input case, an analogous or similar case is retrieved from long-term memory transferred to working memory.
2. **Mapping:** The two cases (the input case and the retrieved analogous case) are “aligned” in terms of their analogous features. This enables the identification of their common and distinctive properties and the inference of unknown properties of the input case based on the properties of the retrieved case.

Clearly, one of the problems of the above procedure is that mapping should already be part of the retrieval process.

Arguably, analogical reasoning is closely related to bisociative reasoning, in particular its domain-crossing conceptual space (long-term memory) bears the hallmarks of bisociation. Furthermore, the concept of “richly interconnected conceptual space in long-term memory” is very similar to the assumption in our formulation of bisociation that there needs to be an overlap of concepts in two domains to facilitate bisociation.

Bisociation is different to analogical reasoning in a number of ways. First, while analogy may be a mechanism in some forms of bisociation, bisociation is not about analogy per se. Perceiving a problem *simultaneously* from the perspective of two distinct knowledge bases, does not mean that one views the entire problem from one knowledge base and then from the other. In a sense, when bisociation occurs, a fraction of both knowledge bases becomes unified into a single knowledge base in the context of the problem at hand. Also, when one considers some of the examples Koestler describes in the context of humor, it is clear that some of these do not rely on the concept of analogy [20]. The

Eureka act described in Figure 2 does not seem to be an example of analogical reasoning. Second, in contrast to bisociation, analogical reasoning seems to suggest a similar (analogous) structure of the long-term memory entity that is retrieved and the input case prompting the retrieval. Bisociation is more akin to Minsky's concept of *knowledge lines* [26], which are a kind of "scaffold" attached to the "mental agencies" (facts, concepts, routines, habits, associations) that were active in creating a certain idea or solving a particular problem in the past. The knowledge lines later work as a way to re-activate the same structures in the context of a new problem. Bisociation could be view in similar terms, except that bisociation explicitly models knowledge lines that cut across knowledge bases embodying domain-specific mental agencies. Thus, when bisociation occurs, mental agencies usually (habitually) active in the context of a specific domain, are activated together with mental agencies usually active in another domain. There are also other perhaps more subtle difference between analogical reasoning and bisociation that are not discussed here.

Swanson's Theory. *Swanson's theory* [33], also known as to as "Swanson linking", is based on the assumption that new knowledge and insight may be discovered by connecting knowledge sources which are thought to be previously unrelated. By "unrelated" Swanson originally meant that there is no co-authorship, no citation and no officially stated relationship among the considered knowledge sources. Swanson coined the phrase "undiscovered public knowledge" to refer to published knowledge that is effectively hidden in disjoint topical domains because researchers working in different domains are unaware of each others' work and scientific publications. He demonstrated his ideas by discovering new relationships in the context of biology and other areas. The field of *literature-related discovery* has emerged from Swanson's work. It aims at discovering new and interesting knowledge by associating two or more concepts described in the literature that have not been linked before [21]. *Conceptual biology* is another line of research in this direction – here the idea is to complement empirical biology by generating testable and falsifiable hypotheses from digital biological information using data mining, text mining and other techniques [4,28]. The methodologies from literature-related discovery and Swanson's theory have already been incorporated in conceptual biology. In combination with systems biology, automatic hypothesis generation is being investigated to facilitate automated modeling and simulation of biological systems [2].

The work by Swanson, literature-related discovery and conceptual biology are related to bisociative information exploration in their attempt to discover information across normally disjoint information spaces. Perhaps one aspect that is strikingly different between the Swanson's approach and bisociation is the notion of unrelatedness and topical disjointedness in Swanson. This assumption separates conceptual spaces on the basis of the originators of knowledge. In our definition of bisociation we do not make this distinction. Nevertheless, the Swanson's theory, while being currently focused on literature as its main source of knowledge, is interesting in the context of bisociation. Further investigations are needed to determine how bisociation and Swanson's approach could complement each other.

Computational Creativity in Science. Computational creativity [7] in art, music and poetry has been around for some time. A recent development is computational creativity applied to the fields of science and engineering. For example, the aim of the Symposium on Computational Approaches to Creativity in Science⁷ (Stanford, US, 2008) was to explore (among other things) (a) the role creativity plays in various scientific areas and how ICT-based tools could contribute to scientific tasks and processes, (b) the nature of creativity in search through a problem space and the representation of the search space and the problem description, (c) the role background knowledge plays in aiding and possibly interfering with creative processes in science, and (d) the interactions among scientists that increase creativity and how computational tools could support these interactions.

There was a wide range of contributions at the Symposium which covered themes such as the design of discovery systems; inter-disciplinary science and communication; abstraction, analogy, classification; spatial transformations and comparisons; conceptual simulation; strategies for searching a problem space; the question of how discovery and creativity differs; knowledge acquisition/refinement approaches and systems; knowledge-based and knowledge management systems, and “knowledge trading zones”; and explanations, models and mechanisms of creative cognition.

Computational creativity in science is a fruitful area and also an area in which large amounts of data, information and knowledge are readily available in computer-readable format. Given the specialization of science on the one hand, and the need for inter-disciplinary science to tackle highly challenging problems on the other hand, it seems that computational creativity in science offers a formidable platform to further investigate biosociative information exploration.

8 Discussion and Conclusion

Computational creativity, in particular computational creativity in non-art applications, is a relatively new computing paradigm [15,8]. For example, computational creativity in science and engineering means that a scientist or an engineer cedes part of her control over the discovery or design process to a computer system that operates with a degree of autonomy, and contributes to the results. In this article we have outlined a rationale or framework for computational creativity based on Koestler's concept of bisociation [20]. The framework presented here facilitates bisociation by “connecting” the knowledge bases of an intelligent agent in the context of a concrete problem, situation or event (Figure 1).

Koestler's treatise and other accounts of bisociation often illustrate bisociation by either bisociating two common or general knowledge domains, or by bisociating one more specialized subject matter domain with a commonsense knowledge domain. For example, the Eureka act (Section 4) bisociates the commonsense domain of taking a bath with the domain of geometry. If we want to reflect this kind of *structure* in a computational creativity solution for non-art

⁷ <http://c11.stanford.edu/symposia/creativity/>

applications, this would mean that we need to develop a knowledge base reflecting the application domain *and* a knowledge base containing commonsense or general knowledge. A commonsense knowledge base contains the knowledge that most people possess and use to make inferences about the ordinary world [24]. Information in a commonsense knowledge base includes things like ontologies of classes and individuals; properties, functions, locations and uses of objects; locations, duration, preconditions and effects of events; human goals and needs; and so on. A commonsense knowledge base must be able to facilitate spatial, event and temporal reasoning. Tasks that require a commonsense knowledge base are considered “AI-complete” or “AI-hard”, meaning that it would require a computer to be as intelligent as people to solve the task.

Another approach to bisociation-based computational creativity would require the bisociating of knowledge bases from different non-commonsense domains, for example, biology and quantum mechanics. Here we have a two-fold challenge:

First, we need to somehow provide some form of interoperability of the involved knowledge bases; this is a topic of active research [10]. Our approach to integrating the concepts from different domains is by creating a *heterogeneous information networks* (called BisoNet in this case) from underlying information sources. The topic of mining of heterogeneous information networks and linked data has become an area of very active research in recent years [18,5].

Second, when the content of bisociated concepts are presented to the user, there may be a considerable problem for the user to recognize potentially useful information from the other domain. For example, a life scientist investigating a detailed mechanism in relation to gene regulation and nuclear receptors may be presented with a scientific article in the field of quantum theory that discusses metric tensors in the context of entanglement entropy. Even if the bisociated article is potentially useful, the life scientist may not be able to “see” the usefulness because he does not have the necessary domain knowledge in field of quantum mechanics.

Another issue – that is shared with all approaches to computational creativity – of the presented framework concerns the assessment of whether or not a discovered item, relationship or bisociation is indeed creative in the sense of being *new*, *surprising* and *valuable* (see Definition 1). This problem is analogous to the issue of determining the degree of interestingness or usefulness⁸ of patterns discovered by means of data mining or machine learning techniques [25]. Sosa and Gero [31] argue that creativity is a social construct based on individual-generative and group-evaluative processes. This suggests that the assessment of creativeness needs to incorporate social aspects that transcend the within-individual cognitive dimension. This points to a rather complex challenge for computational creativity and is something that future studies of computational bisociation need to take on board.

⁸ In addition to these, the discovered patterns are usually also required to be non-trivial, valid, novel and comprehensible. Depending on the technique used and the application area, an automated assessment of these additional dimensions may also pose a considerable challenge.

With the increasing power of ICT and the growing amounts of data, information and knowledge sources, there is a new wave of efforts aiming to construct computing solutions that exhibit creative behavior in the context of challenging applications such as science and engineering [8]. This article presents a framework for computational creativity based on the concept of bisociation [20]. As a pioneering effort in this field, the BISON project⁹ has been exploring bisociation networks for creative information discovery. This article presents some of the rationale, ideas and concepts we have explored in an effort to formally define the concept of bisociation and bisociative information exploration. Clearly, more work is needed to develop a more comprehensive formal understanding of bisociation and how this concept can be used to create novel ICT methods and tools.

Open Access. This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Adams, M.P.: Empirical evidence and the knowledge-that/knowledge-how distinction. *Synthese* 170, 97–114 (2009)
2. Ananiadou, S., Kell, D.B., Tsujii, J.-I.: Text mining and its potential applications in systems biology. *Trends in biotechnology* 24(12), 571–579 (2006)
3. Barron, F.: Putting creativity to work. In: Sternberg, R.J. (ed.) *The Nature of Creativity*, pp. 76–98. Cambridge University Press, Cambridge (1988)
4. Bekhuis, T.: Conceptual biology, hypothesis discovery, and text mining: Swanson's legacy. *Biomedical Digital Libraries* 3, 2 (January 2006)
5. Bizer, C., Heath, T., Berners-Lee, T.: Linked data – the story so far. *International Journal on Semantic Web and Information Systems (IJSWIS)* 5(3), 1–22 (2009)
6. Boden, M.A.: Précis of the creative mind: Myths and mechanisms. *Behavioural and Brain Sciences* 17, 519–570 (1994)
7. Boden, M.A.: Computer models of creativity. In: Sternberg, R.J. (ed.) *Handbook of Creativity*, pp. 351–372. Cambridge University Press, Cambridge (1999)
8. Bridewell, W., Langley, P.: Symposium on computational approaches to creativity in science. Final report for NSF grant IIS-0819656. Technical report, Institute for the Study of Learning and Expertise (2008), <http://c11.stanford.edu/symposia/creativity/text/scacs08.report.pdf>
9. Chalmers, D.J., French, R.M., Hofstadter, D.R.: High-level perception, representation, and analogy: A critique of artificial intelligence methodology. *Journal of Experimental and Theoretical Artificial Intelligence* 4, 185–211 (1992)
10. Chaudhri, V.K., Farquhar, A., Fikes, R., Karp, P.D., Rice, J.: OKBC: A programmatic foundation for knowledge base interoperability. In: *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI 1998)*, Madison, Wisconsin, USA, pp. 600–607 (1998)
11. Einstein, A.: Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik* 17, 132–148 (1905)

⁹ Website of the BISON (Bisociation Networks for Creative Information Discovery) project: <http://www.bisonet.eu>.

12. Falkenhainer, B., Forbus, K.D., Gentner, D.: The structure-mapping engine. *Artificial Intelligence* 41, 1–63 (1989)
13. Gentner, D.: Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7(2), 155–170 (1983)
14. Gentner, D., Colhoun, J.: Analogical process of human thinking and learning. In: Glatzeder, B.M., Goel, V., von Muller, A. (eds.) *Towards a Theory of Thinking. On Thinking*, pp. 35–48. Springer, Heidelberg (2010)
15. Gero, J.S.: Computational models of innovative and creative design processes. *Technological Forecasting and Social Change* 64(2-3), 183–196 (2000)
16. Good, J.: A five-year plan for automatic chess. *Machine Intelligence* 2, 110–115 (1968)
17. Hall, R.P.: Computational approaches to analogical reasoning: A comparative analysis. *Artificial Intelligence* 39, 39–120 (1989)
18. Han, J.: Mining Heterogeneous Information Networks by Exploring the Power of Links. In: Gama, J., Costa, V.S., Jorge, A.M., Brazdil, P.B. (eds.) *DS 2009. LNCS (LNAI)*, vol. 5808, pp. 13–30. Springer, Heidelberg (2009)
19. Higgins, J.M.: 101 Creative problem solving techniques. New Management Publishing Company (1994)
20. Koestler, A.: *The act of creation*. Penguin Books, New York (1964)
21. Kostoff, R.N.: Literature-Related Discovery (LRD): Introduction and background. *Technological Forecasting and Social Change* 75(2), 165–185 (2008)
22. Kumar, M.: *Quantum: Einstein, Bohr and the great debate about the nature of reality*. Icon Books Ltd. (2008)
23. Kurzweil, R.: *The age of intelligent machines*. Massachusetts Institute of Technology (1990)
24. Lenat, D.B., Guha, R.V., Pittman, K., Pratt, D., Shepherd, M.: *Cyc: Toward programs with common sense*. *Communications of the ACM* 33, 30–49 (1990)
25. McGarry, K.: A survey of interestingness measures for knowledge discovery. *The Knowledge Engineering Review* 20(1), 39–61 (2005)
26. Minsky, M.: K-lines: A theory of memory. *Cognitive Science* 4, 117–133 (1980)
27. Morrison, C.T., Dietrich, E.: Structure-mapping vs. high-level perception: The mistaken fight over the explanation of analogy. In: *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, pp. 678–682 (1995)
28. Natarajan, J., Berrar, D., Hack, C.J., Dubitzky, W.: Knowledge discovery in biology and biotechnology texts: A review of techniques. *Critical Reviews in Biotechnology* 25(1/2), 31–52 (2005)
29. Schaeffer, J., Culberson, J., Treloar, N., Knight, B., Lu, P., Szafron, D.: A world championship caliber checkers program. *Artificial Intelligence* 53(2-3), 273–290–115 (1992)
30. Smith, E.E., Medin, D.L.: *Categories and concepts*. Harvard University Press, Cambridge (1981)
31. Sosa, R., Gero, J.S.: A computational framework for the study of creativity and innovation in design: Effects of social ties. In: Gero, J.S. (ed.) *Design Computing and Cognition 2004*, pp. 499–517. Kluwer Academic Publishers, Dordrecht (2004)
32. Sowa, J., Majumdar, A.K.: Analogical Reasoning. In: Ganter, B., de Moor, A., Lex, W. (eds.) *ICCS 2003. LNCS (LNAI)*, vol. 2746, pp. 16–36. Springer, Heidelberg (2003)
33. Swanson, D.R.: Fish oil, Raynaud’s syndrome, and undiscovered public knowledge. *Perspectives in Biology and Medicine* 30(1), 7–18 (1986)