

The Input, Coherence, Generativity (ICG) Factors. Towards a Model of Cognitive Informativity Measures for Productive Cognitive Systems.

Ana-Maria Oltețeanu

Cognitive Systems, Universität Bremen
Bremen, Germany
amoodu@informatik.uni-bremen.de

Abstract. Classical thinking on information and informativity considers the *informee* as a perfect information receiver. However, when studying productive natural and artificial cognitive systems, cognitively based models of informativity need to be formulated. Three factors relevant to cognitive informativity measures are proposed: Input, Coherence and Generativity (ICG). These factors take into account the type of Input which can be stored, the Coherence of the system after acquiring the information, and the Generativity of the system after the new information was integrated.

Keywords: Cognitive Systems, Cognitive Informativity Measures, Creative systems, Productive systems, Generativity, Coherence

1 Introduction

Various ways of thinking about information exist [29, 7], which imply different ways of measuring informativity. Between information as data communication, and information as semantic content, informativity either considers data in its probabilistic nature and ability to surprise the *informee*¹[28], or its well-formedness, meaningfulness or truthfulness.

However, such informativity measures seem to consider the *informee* as a perfect information receiver - one that can comprehend, memorize and utilize whatever information it is given. This does not cover the perspective of productive natural or artificial cognitive systems, which enter any information gathering act with cognitive economy priorities, computational constraints, and which dynamically use information and their knowledge in order to produce more knowledge and create new artifacts.

The measures of cognitive informativity proposed here stem from general cognitive principles of information representation, structuring and processing. They reflect the subjective nature of every informational act as involving the

¹ The information receiver.

interaction between the *informant*² and an individual informee. Any individual informee is seen as a cognitive system. Far from being an empty vat, an informee already holds various types of knowledge structures, which can be used for (or stand in the way of) the acquisition of new knowledge (the Input factor).

Furthermore, the consequences of the informant being integrated in the knowledge base (KB) of the informee can vary, depending on whether such added informant makes the knowledge base more coherent, easier to navigate, or adds to its entropy (the Coherence factor).

Finally, as we consider the capacity of a cognitive system to be productive (i.e. able to creatively solve problems, create new artifacts, make new inferences) to be essential, the third factor in this cognitive informativity model relates to whether the informant leaves the informee with a better or worse ability to generate new cognitive artifacts (the Generativity factor).

The main contribution of this paper is thus to lay the foundations for a model of cognitive informativity under the general paradigm of productive cognitive systems.

Ways to evaluate computational creativity systems have been proposed [32, 26, 25, 4], however they deal with the assessment of an artefact or a process as creative, not with the impact an informant has on the generativity of a system. Steps towards the cognitive modeling of creativity have been made [30, 22, 23], some of which propose approaches in which structure is relevant. However, the impact a new informant has on the knowledge base of the informee has not been assessed from the perspective of cognitive creative systems.

The rest of this paper is structured as follows. First, arguments from cognitive science and AI which support the direction of the cognitive informativity measures further proposed are presented in Section 2. Based on these arguments, an initial analytic toolkit is laid down in Section 3.

Taking into account the information already present in the informant's knowledge base and its structure, three main factors relevant to setting cognitive informativity measures are described - Input, Coherence and Generativity. Each of these factors is then analysed in turn, together with ensuing possible informativity measures (Sections 4-6). Implications of the model are discussed and further work proposed in Section 7.

2 Setting the stage

To further clarify why cognitive informativity measures are necessary and set the analytic toolkit for defining important factors, in this section: (1) the case for structured representations in AI and cognitive science is discussed; (2) primary tools for estimating measures of informativity are presented and (3) the diversity of ways of knowing available to a cognitive system is discussed and exemplified.

² The information transmitted.

2.1 The case for structured representation

Structured representations are posited quite often in cognitive science and artificial intelligence. Depending on (i) the task the system needs to solve (in AI), or (ii) the task the system solves in a manner which needs understanding and explaining (in cognitive science), such representations range from image schemas [14], frames [16], scripts [27], to spatial templates, conceptual spaces [8] and mental models [12]. All such terms make the case for structured representation, which has implementation, modeling and comprehension value in cognitive systems. They also point towards the need for structure and coherence in a cognitive agent's knowledge base.

This structure might be an empirical reality of human minds, or simply an artefact of the quest for modeling cognitive systems (a debate which won't be engaged here). However, analysing the interaction between such representational structure and information that enters the system can be useful in formulating cognitive measures of informativity.

2.2 Using representation structure for measures of informativity

Natural cognitive systems generally have knowledge bases which are dynamic and plastic [24, 3]. This supports their ability to learn, to recall and change information previously held. However, acting on and manipulating this knowledge base, which can be thought of as long-term memory, encounters the bottleneck of working memory [1]. Thus research on cognitive load [31] and processes of productive/creative thought like analogy [9, 11] and metaphor [13] indicate that new information is often compared to information previously held in the knowledge base, and the ability to ground previous information in older structures generally helps comprehension, the knowledge acquiring process and the production of new information.

To add new information to the system, the nature and structure of the previously held information is thus relevant. In order to propose cognitive measures of informativity which can define the impact a piece of information can have on a cognitive system, metrics which reflect the structure and elements of information held in the initial knowledge base of the system are needed.

No normative comments on what such representation structures should be are made here. However, a theoretical framework with hybrid representations which can be used to further define informativity measures has been described [20, 21] and will be used as a further tool for the following examples. We summarize some of the principles of this framework in the following.

Take a cognitive system A , with a set of sensory modalities S and $S_1, S_2, \dots, S_n \in S$. The features encoded through such sensory modalities are categorized in feature maps, depending on the sensory modality they elicit. Such feature maps are then used for comparison whenever encoding or recognizing a new but similar object, and help future creative problem-solving endeavours of the cognitive system [20]. Each object seen is an activation of sensory features in these respective feature maps. These maps can contain knowledge of color, shape, motion and

any other thing which can be gathered via sensory input, including information on the state of the agent. Objects and other concepts are further encoded as representation structures RS [20]. From the knowledge representation perspective, both concrete and abstract concepts are collections of data (activation of features spaces) which cluster in a viable RS (viability is defined below). Abstract concepts and more complex representation structures are anchored on other concrete concepts. Furthermore, these can be assembled in meaningful groups of RS based on context (what objects are encountered together), consequence (what happens if a particular motion is initiated with specific types of objects) and include interpretations for the agent (what the consequences of particular strings of events or spatial positioning of objects, etc. will be on the goals and needs of the agent).

From this perspective, a new *informant* can be:

- (i) a new concept - which will normally create a new RS ;
- (ii) information on a previous concept - which adds to a previous RS or in other ways rectifies already held information;
- (iii) a new higher level RS which is added based on already held concept(s) - e.g. A learns to use previously known objects/concepts in the context of a new problem;
- (iv) a higher level RS is modified - e.g. A learns new consequences for a previously held routine, or learns to integrate a new object in that routine, etc.

Learning new concepts or higher level representations structures like in (i) and (iii) is a costly process, and various prerequisites might need to exist for this, like the concepts for encoding the higher RS (see Section 4). Modifications of previous RS like in (ii) and (iv) might depend: on a) the strength of encoding of previous information or b) the strength (measured as importance or salience) of the new information shown. Generally, the pre-existing structure of the KB might play a determining role on whether the *informant* can be integrated (see Section 5).

This framework of representation is general enough to allow further adaptation and thus support the discussion about informativity measures in various systems, cognitive architectures and the human mind. To explain the proposed cognitive informativity measures, representation structure (RS) will be used as a base unit (with the plural RS s), where a viable RS is a stable, meaningful, useful set of features in the knowledge base of the system or agent.

These viability constraints to representation structures stand for a bare minimum, and might be incomplete. This is a short description of each:

- A RS generally has to be stable - it can be encoded with ease by the system, no strong competition exists in parsing the features in the subsymbolic plane in a different RS . This is a noteworthy point as various interpretations of the same sets of features can exist, as illustrated by bistable perception [15]. When the representation structures build on pre-existing concepts, differences in encoding at lower levels can propagate upwards in abstract concepts or higher level representation structures. Furthermore, higher RS might be encoded in different sets of concepts from agent to agent. However, stable representation seems to be the

norm, possibly based on neural and cognitive economy grounds, with changes in representation being the event.

- A *RS* has to be meaningful - representation structures are place holders for something in or about the world (which includes the agent itself, its interaction with the world and internal world), or about other *RSs*. The agent uses *RSs* to make sense of the world. Their degree of accuracy does not necessarily overlap with their meaningfulness, as such placeholders are needed by the cognitive system for a world interpretation³. Some *RS* can be more meaningful (and thus more influential) than others, gathering together an interpretation of the world and the position of the agent in it, or holding a keystone role within the system.

- A *RS* has to be useful. This reflects the overt or covert interest of the agent in acquiring or creating that *RS*. Such usefulness can be reflected at many levels: thinking about the world and understanding it (meaningfulness), solving problems (functional), creating new artifacts (expressive), some of which solve problems (expressive/functional).

2.3 Ways of knowing

Various ways of knowing an object or concept can be accounted for by: (1) different sensory experiences, (2) different types of knowledge representations, (3) different associative links and (4) relations encoded for those representations. For example:

(1) Agent A_1 knows object O_1 with sensors S_1, S_2 , while agent A_2 knows object O_1 with sensors S_2, S_3 .

For example, let's take $O_1(KB_{A_X})$ to represent the knowledge agent A_X has in its KB about object O_1 . If $A_1 = John, A_2 = Mary, O_1 = curtain, S_1 = size, S_2 = color, S_3 = texture$:

- $curtain(KB_{John}) \supset \{200cm*220cm, orange_5^4\}$ - John has seen the curtain in someone's house but didn't touch it.

- while $curtain(KB_{Mary}) \supset \{orange_6, soft_3\}$ - Mary has seen and touched a sample of the curtain in a catalogue;

(2) Even if sensor S_2 is shared, feature $f_{(S_2, O_1)}$ can be categorized in KB_{A_1} and KB_{A_2} in proximity to different features, depending on what else each agent already knows/stores.

Thus, in our example, the orange perceived by John might have been classified close to *bright orange* or the color of another concept - *fire*, while Mary might classify it as *Salmon orange* or next to *pink*;

³ If no other *RSs* are known, or the cognitive cost of replacing them is too high, possibly by having a knock-on effect of destroying (entire systems of) other *RSs*, *RSs* might be kept in place even when proven wrong by the environment. Errors of judgement and biases are common place for natural cognitive systems. This points to meaningfulness being an important requirement (and possibly more important than truthfulness.)

⁴ Shades of color and texture are represented here with numbers as to reflect their perception by a visual or tactile sensor.

(3) Encoding of $f_{(S_2, O_1)}$ in different places in the sensory map, and in different object contexts yields different associative links in KB_{A_1} and KB_{A_2} .

Thus for John this might yield an association to *sunset*, while to Mary an association to *candy*;

(4) Such different encoding can predispose future ease of access, making $f_{(S_2, O_1)}$ easier to interpret by different relations, out of which its embedding in further *RSs* might depend.

Thus John might want the curtain for his bedroom, while Mary might think through her association to candy that it is only appropriate for a child's room. Knowing size might enable John to come up with the creative use "*can be used as a rope*" for the curtain. While knowing texture, Mary could come up, for the same curtain, with the creative uses "*can be used wrapped as a pillow, or to clear leaks if nothing else is in hand*".

Cognitive systems are dynamic systems which can represent various properties of the objects at hand, depending on their own goals, interest, context, or recent biases. Conversely, various ways of knowing a property can be anchored in knowledge about different sets of objects or more concrete concepts which contain that property. In a sense, knowledge about any object is always incomplete - as further knowledge could always be collected if one would possess a different type of sensor, would have studied the object from a different angle or would know about other significant functional relations of the object to other objects. This makes human communication imprecise, however it doesn't impede it. It just emphasizes the fact that two agents that both "know" the same object, might hold overlapping but different types of information and context about it.

From the perspective of productive cognitive systems, this imperfection in communication has problem-solving advantages. Thus, *informants* transmitted by A_1 to A_2 might not be considered in the same knowledge configuration, but might fit *RSs* held in KB_{A_2} and help solve a problem P_x for A_2 , which they wouldn't have solved for A_1 due to the way the informant was embedded in KB_{A_1} or its representation structures.

One could retort that, as the knowledge in artificial cognitive systems can be (in many encoding paradigms) examined, two artificial agents can indeed bring the same knowledge, representation structure or interpretation of an external object to the table, and that in this case, the grounding of the object is perfectly equivalent (the agents are both talking about exactly the same thing). However, if these systems apply cognitive processes which include an ability to comprehend knowledge in relation to other knowledge held in their knowledge base, this equivalence depends on how each object is positioned within the *KB* itself, what other knowledge is closely related, what other objects are similar, etc. Two artificial systems would need to have exactly the same knowledge in the same knowledge base structure to be equivalent. This could very well constitute a case 0 of perfect communication (requiring 0-change for comprehension), in which cognitive informativity measures would be irrelevant, and receiving an *informant* would have an equivalent impact on both systems (due to their similar knowledge structure).

3 Principles of organization for a cognitive account of informativity measures

Some subjective matters relevant for any information-transfer accounts are left aside in the following.⁵ The account of informativity proposed here is based on a comparison between the *informant* and knowledge already present in the cognitive system, its organization and productive capacity.

In the following, three factors that such an account needs to address are proposed and discussed. These three factors are:

1. The ways in which the information can be added to the system;
2. The coherence in the system;
3. The generativity of the system after the information addition.

Adding information to a cognitive system presupposes the system's ability (and sometimes interest) in encoding it. Here, the subjective matter of the system's interest is not addressed, only its ability and the effort required to encode the *informant* based on its sensors, the grounding hypothesis and organization of the knowledge it already holds.

The coherence in the system's knowledge base involves both lack of contradictions and the ability to integrate and connect new and old information. This is relevant to the system's further ability to utilize the *informant*, or the parts of the informant it has managed to integrate. This factor points to general measures of informativity which assess the impact the informant has on the structure and connectivity of the knowledge base.

The generativity of the system is its ability to be productive (of artifacts, solutions, ideas). Under the generativity factor, a measure of informativity is proposed that reflects the change in productive capacity of a system after it has received a particular information.

4 Adding information to a cognitive system (Input)

Developing on previous accounts of the grounding problem [10] and hypotheses of how such grounding can be solved [2], in order for a cognitive system to be able to represent, understand, memorize and utilize information, such a system needs to be able to ground this information in sensory modalities or in previously acquired knowledge. This points to two main question clusters about the informant which are relevant from the cognitive system's perspective:

1. Is the information groundable in the system, or representable by different information structures? Does the information fit pre-owned structures? Do the elements necessary for encoding the new piece of information exist? With how much accuracy can the new information be encoded by the *RSs* in the knowledge base? (the stable anchoring question)

⁵ These include, but are not limited to: goals of the system that receives such information, its previous biases, current mood - an exploration-learning wide-focus versus exploitation narrow-focus, etc.

2. Is the information a new *RS*, does it present features of previous *RSs*, or does it built upon them at a higher, more abstract level? (the type of informant with raport to knowledge organization question)

The first question asks whether elements required for representation are present (anchoring perspective). The more elements present, the easier the information will be represented and therefore memorized and comprehended. The easier the information is represented, the higher level *RSs* can be built on top of it. The first question also refers to whether competitive *RSs* might exists, which might prevent such easy representation and grounding. If the informant is too close to something the system already holds in its knowledge base, or provides a concurrent but different representation or interpretation, this might restrict access to the previous knowledge, or it might prevent categorization.

The second question addresses where the informant can be embedded in the knowledge base of the system. In this context, effort of such embedding can be discussed. Such effort has to take into account: (i) sparsity of data; (ii) competition in categorization in the knowledge base in which the encoding is done (after the previous question has decided whether the encoding is at all possible), and (iii) the amount of data that needs encoding. Thus acquiring an entire new *RS* versus acquiring a feature for a previous *RS* might be comparable, from the information-processing perspective, to the costs of adding a new class or data-structure in an already working program, versus adding a new instantiation of a former class. However, due to specific categorization-competition constraints in cognitive systems, encoding an entirely new *RS* might be easier than encoding and using an *RS* which competes with already encoded *RSs*, because of the reorganization necessary in the knowledge base to accomodate the new informational refinement.

A non-exhaustive list of feature-*RS* and *RS-RS* fitting examples is presented in the following:

- (i) new feature f_y fits previous unfitted but present slot in RS_x
- (ii) new feature f_y can be added in a coherent non-competitive manner to RS_x
- (iii) new feature f_y fits RS_x but as a consequence RS_x has to change;
- (iv) new feature f_y competes with already encoded feature f_z in R_x ;
- (v) new RS_y can be anchored in elements already present in the KB (e.g. fits f_a, f_b, f_c), without competition;
- (vi) new RS_y can be anchored in elements already present in the KB (e.g. fits f_a, f_b, f_c), and is necessary for grouping these element (compression constraints) or solving a problem;
- (vii) new RS_y can be anchored in elements already present in the KB , but will encounter competition from RS_z which is a different interpretation or category anchored in overlapping elements;
- (viii) new RS_y fits higher level RS_z , etc.

5 The Coherence in the system

Coherence in the system can be a measure of both the system's lack of internal contradictions, as well as a measure of the system's general connectivity. The

latter is addressed in the following, because of its relevance to productive systems. Coherence (as connectivity) in a system is responsible for easy navigation or flow between encoded pieces of knowledge. This impacts the system's ability to search for information, manipulate it and change it in a productive manner.

Let's say a system uses links between its representation structures to navigate its knowledge base. Various types of such links can be envisaged: (i) associative links (RS_a is like RS_b in some way), (ii) relational links (which put two representations in a relation to each other), (iii) functional links (which help express a RS through a set of concepts or other RSs), etc.

Whether the coherence of the system is increased or decreased with the addition of new information should be an important factor in assessing an *informant's* impact on a cognitive system. Confirmation bias [18] and cognitive dissonance [6] show that cognitive systems generally aim to keep a high level of coherence. Attaining a higher level of coherence can make the system more productive: a RS_x which associates with an older RS_a so that a $NewLink(RS_x, RS_a)$ is produced, might provide new avenues to navigate the knowledge base, while a RS_y which produces a new association between two previous representation structures RS_a and RS_b , so that $NewLink(RS_a, RS_y, RS_b)$, will increase not just coherence but might make the system able to construct a new representation structure $RS_z = NewLink(RS_a, RS_y, RS_b)$.

Informativity measures of a particular informant might be computed here in terms of the state of the associative, relational or functional links in the system after the informant has been integrated. Take the number of such links in the KB of agent X at time t_y to be represented by $NoOfLinks(KB_X, t_y)$; the size of a particular RS_z at time t_y to be represented by $LengthOf(RS_z, t_y)$ and the number of unconnected RSs in KB_X at time t_y to be $NoOfUnconnectedRSs(KB_X, t_y)$. An informant providing higher coherence could have as a consequence:

$$NoOfLinks(KB_A, t_1) \geq NoOfLinks(KB_A, t_0);$$

$$LengthOf(RS_z, t_1 | RS_z \in KB_A) \geq LengthOf(RS_z, t_0 | RS_z \in KB_A)^6;$$

$$NoOfUnconnectedRSs(KB_A, t_1) \leq NoOfUnconnectedRSs(KB_A, t_0);$$

when t_0 is the previous and t_1 the consequent time state at which the coherence of the KB is assessed.

However, coherence decrease might also be beneficial long-term, even if it increases the short-term entropy of the system, or it leaves the system with too many open questions. Knowing pieces of information in other fields might open the system for future grounding of novel information, while a system that is completely coherent might also be closed due to its stability. However the effort of encoding and holding in memory unconnected pieces of information might initially be greater, despite their possible further uses. Cognitive systems could further be defined through the individual threshold they put on lack of coherence.

New information can change the structure of the system's knowledge base (and its emphasis). Thus an important question in terms of knowledge base integration is whether a new informant: (a) adds to existing structure, (b) helps

⁶ Note that this does not even start to tackle the differences in structure between RS_z at time t_1 and RS_z at time t_0 .

connect it further or (c) helps parse it in any different way. Information that helps re-representing or re-interpreting older information (adds to the ways of restructuring) would be particularly important from the generativity factor point of view of productive systems.

6 The Generativity in the system

As mentioned earlier, information that helps future restructuring [5, 19] can be useful for future dynamic progress of the system, and for its productivity - for its ability to see objects and events of the world in a different interpretative light, and come up with new representation structures that solve problems or otherwise generate useful artifacts.

Generally speaking, generativity is a matter of how much more new information can be produced by the system. This can be taken to mean how many new information-combination structures (that are stable) are possible, but also what possibilities of discovery of new unknown elements are made available through bringing forth more flexibility in restructuring the old *KB* in new ways.

Thus new combinations can be achieved through transfer of previous *RSs* to new feature sets, or learning new *RSs* which can interpret previous feature sets. Others are combinations of previous *RSs*. Such processes might be spontaneous or serendipitous, they might be due to cognitive and/or environmental context, emergent due to overlapping features, or belabored. Some new combinations are not complete - they point to higher-arching *RSs* that miss elements, thus revealing gaps in knowledge.

One could consider all such possible new structures to represent the generativity capital of the system. However, in reality, due to cognitive economy principles, such transfer is hard without previous connectivity in the system. A *RS* will rarely be transferred to completely new features, unless it fits them in an unusually stable way. Two *RSs* might rarely connect without some previous common features or some overarching new *RS* that comes as a response to a problem. The Coherence factor might thus have instrumental consequences on the Generativity factor.

In conclusion, the impact of a new informant on the generativity of a productive cognitive system can be assessed using questions of the following type: given the original structure of the system, does the informant contain new features or *RSs* such that:

- new (stable, meaningful, useful) *RSs* or external artifacts can be produced?⁷
- new connections are possible?
- new elements might be revealed as unknown?
- new restructuring possibilities appear? etc.

⁷ A cluster of connections might not yield a full blown representation structure, but with the help of external expression and cognition, they might be compelling enough for the cognitive system to explore, thus leading to new productive processes.

7 Discussion and further work

After describing the Input, Coherence and Generativity factors, a tentative definition of cognitive informativity, based on the relation between the information transmitted and information receiver, can be made. Thus cognitive informativity as defined so far is a measure of (i) the effort required to ground and integrate the new information, (ii) the changes that occur in the knowledge structure of the informant and (iii) the influence this information has on the productivity of the system.

One might ask what the purpose of such informativity measures might be, considering that the knowledge base of natural cognitive systems could be very hard to estimate.

First, as we work towards artificial cognitive systems that can at least mimic if not implement some of the adaptive powers of their natural counterparts, it is important to remember that the knowledge in such systems can be measured (and so is coherence, generativity) - therefore such cognitive informativity measures can be applied to artificial systems. For example, in the case of OROC [23], adding knowledge about a new object to the KB might provide further generativity to the system. As OROC uses knowledge about similar objects to infer affordances, if an informant concept c_4 (*a flowerpot*) is of a similar shape and material as a known object c_3 (*a cup*), but has an affordance which is unknown to the system (*to grow flowers in*), OROC will make the creative inference that the older known object c_3 might be used in the same manner, coming up with the creative use "*Maybe we can use a cup to grow flowers in.*". Thus OROC's generativity will increase when it receives informants that can be coherently integrated within its KB. In OROC, such changes in generativity and coherence could be measured.

Second, an open discussion of the limits implied by imperfect informativity might bring further analytic tools for the learning sciences, processes of communication, communication theory and HCI. A first estimate which compares the initial knowledge base of the informee to the informant can yield further assessments of the ways in which the informant should be communicated, or previous representational structures can be put in place, thus smoothing the communication or learning process, making clearer the possible gaps of knowledge and significantly increasing the informativity of the informant.

Third, if the information processing metaphor is taken quite far into the cognitive science ground, the principles of cognitive economy and measures of computational complexity could imperfectly align to be each other's counterparts. This means that one should think about computational complexity under a bounded rationality paradigm, in terms of generativity (what avenues can a cognitive system explore), rather than perfect variants (given enough memory and time, can a system get there). This perspective centers on the informativity a system can obtain from its environment given its internal structure (thus it is individualized and goal-oriented). Artificial agents with limited knowledge and defined knowledge structures, processing in real-time noisy environments could strongly benefit from it.

The purpose of this paper has been to introduce the idea of a need for cognitive measures of informativity for productive systems, and to propose a model based on three factors which affect cognitive systems - Input (grounding), Coherence and Generativity. This introduction can benefit from further work which will formalize in a rigorous manner the measures of informativity proposed here as a function of each factor.

Acknowledgments. The support of the Cognitive Systems group and of the Impulse program at the Universität Bremen is acknowledged. We thank the reviewers for their constructive suggestions, which helped improve this paper.

References

1. Baddeley, A. D., & Hitch, G.: Working memory. *Psychology of learning and motivation*, 8, 47-89 (1974).
2. Barsalou, L.: Abstraction in perceptual symbol systems. In *Philosophical transactions of the Royal Society of London*, volume 358 of B, Biological sciences, pages 1177-87 (2003)
3. Caroni, P., Donato, F., & Muller, D.: Structural plasticity upon learning: regulation and functions. *Nature Reviews Neuroscience*, 13(7), 478-490 (2012)
4. Colton, S., Pease, A., & Charnley, J.: Computational creativity theory: The FACE and IDEA descriptive models. In *Proceedings of the Second International Conference on Computational Creativity* pp. 90-95 (2011).
5. Duncker, K.: On problem solving, *Psychological Monographs*, no. 5, whole no.270, vol. 58, (1945)
6. Festinger, L.: Cognitive dissonance. *Scientific American* 207 (4): 93-107 (1962)
7. Floridi, Luciano: The method of levels of abstraction, *Minds and Machines*, no.3, vol. 18, 303-329, Springer (2008)
8. Gärdenfors, P.: *Conceptual Spaces: The Geometry of Thought*. Bradford Books (2004)
9. Gentner, Dedre: Structure-mapping: A theoretical framework for analogy. *Cognitive science*, no. 2, vol. 7, 155-170, Elsevier (1983)
10. Harnad, Stevan: The symbol grounding problem. *Physica D: Nonlinear Phenomena*, no. 1, vol. 42, 335-346, Elsevier (1990)
11. Holyoak, K. J.: *Mental leaps: Analogy in creative thought*. MIT press (1996).
12. Johnson-Laird, P.N., *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*, Cognitive science series, Harvard University Press (1983)
13. Lakoff, George and Johnson, Mark: *Metaphors we live by*, University of Chicago Press, Chicago (1980)
14. Lakoff, George and Johnson, Mark: *Philosophy in the flesh: the embodied mind and its challenge to Western thought*, Basic Books, New York (1999)
15. Mamassian, P., & Goutcher, R.: Temporal dynamics in bistable perception. *Journal of Vision*, 5(4), 7 (2005)
16. Minsky, M.: A framework for representing knowledge. In Winston, P., editor, *The Psychology of Computer Vision*, pages 211-277. New York: McGraw-Hill (1975)
17. Newell, A. and Simon, A.: *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall (1972)

18. Nickerson, R. S.: Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, 2(2), 175 (1998)
19. Ohlsson, S.: Information processing explanations of insight and related phenomenon. In *Advances in the Psychology of Thinking*, eds. Keane M. and Gilhooly K., 1–44, London: Harvester-Wheatsheaf (1992)
20. Olteţeanu, Ana-Maria: Two general classes in creative problem-solving? An account based on the cognitive processes involved in the problem structure - representation structure relationship. In *Proceedings of the Workshop “Computational Creativity, Concept Invention, and General Intelligence”*, eds. Besold, T. ; Kühnberger, K.-U.; Schorlemmer, M. and Smaill, A., Publications of the Institute of Cognitive Science, 01-2014, Osnabrück (2014)
21. Olteţeanu, Ana-Maria: From Simple Machines to Eureka in Four Not-So-Easy Steps. Towards Creative Visuospatial Intelligence. In V. C. Müller (Ed.), *Fundamental Issues of Artificial Intelligence (Synthese Library)*. Berlin: Springer. (to appear)
22. Olteţeanu, Ana-Maria and Falomir, Zoe: comRAT-C - A Computational Compound Remote Associates Test Solver based on Language Data and its Comparison to Human Performance. In “Cognitive Systems for Knowledge Discovery”, ed. Lledó Museros, Núria Agell, and Oriol Pujol, *Pattern Recognition Letters* (to appear)
23. Olteţeanu, Ana-Maria and Falomir, Zoe: Object Replacement and Object Composition in a Creative Cognitive System. Towards a Computational Solver of the Alternative Uses Test. In “From Human to Artificial Cognition (and back): New Perspectives on Cognitively Inspired AI Systems”, ed. Antonio Lieto and Daniele P. Radicioni, *Cognitive Systems Research* (submitted)
24. Pascual-Leone, A., Amedi, A., Fregni, F., & Merabet, L. B.: The plastic human brain cortex. *Annu. Rev. Neurosci.*, 28, 377–401 (2005)
25. Pease, Alison and Winterstein, Daniel and Colton, Simon: Evaluating machine creativity. In *Workshop on Creative Systems, 4th International Conference on Case Based Reasoning*, 129–137 (2001)
26. Ritchie, Graeme: Assessing creativity, in: Wiggins, G.A. (Ed.), *Proceedings of the AISB’01 Symposium on AI and Creativity in Arts and Science* (2001)
27. Schank, R. C. and Abelson, R. P.: *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum (1977)
28. Shannon, C.E. and Weaver, W.: *The mathematical theory of communication*. Urbana: University of Illinois Press (1949)
29. Shannon, C.E.: *Collected papers*. Ed. Sloane, N.J.A. and Wyner, A.D., IEEE Information Theory Society, New York: IEEE Press (1993)
30. Hélie, Sébastien and Sun, Ron: Incubation, insight, and creative problem solving: A unified theory and a connectionist model. *Psychological review*, 117 (3). 994, American Psychological Association (2010)
31. Sweller, J.: “Cognitive load during problem solving: Effects on learning”. *Cognitive Science* 12 (2): 257–285 (June 1988).
32. Wiggins, Geraint A: Towards a more precise characterisation of creativity in AI, in *Case-based reasoning: Papers from the workshop programme at ICCBR, vol. 1*, 113–120 (2001)