Formal Ontologies and Semantic Technologies: A "Dual Process" Proposal for Concept Representation

Marcello Frixione[1], **Antonio Lieto**[2]

[1] DAFIST, University of Genova, via Balbi, 4 - 16126 Genova, Italy
[2] Dept. of Computer Science, University of Torino, C.so Svizzera, 185 - 10149 Torino, Italy
and ICAR-CNR, Viale delle Scienze Ed. 11, 90128 Palermo, Italy
[1]*marcello.frixione@gmail.com*, [2]*lieto@di.unito.it*

PREPRINT of the penultimate version submitted, the final version will be available on the forthcoming issue of *Philosophia Scientiae Journal*.

Abstract One of the main problems of most contemporary concept oriented knowledge representation systems consists in the fact that, for technical convenience, the representation of knowledge in prototypical terms, and the possibility of exploiting forms of typicality-based conceptual reasoning, is not permitted. Conversely, in cognitive sciences, evidence exists in favor of prototypical concepts and non-monotonic forms of conceptual reasoning have been extensively studied. This “cognitive” representational and reasoning gap constitutes a problem for the computational systems, since prototypical information plays a crucial role in many relevant tasks. Taking inspiration from the so-called dual process theories of reasoning and rationality, we propose that conceptual representation in computational systems should be demanded to (at least) two representational components, each specialized in dealing with different kinds of reasoning processes. In this article, the theoretical and computational advantages of such “dual process” proposal are presented and briefly compared with other logic-oriented solutions adopted for facing the same problem.

Abstract L'un des problèmes principaux de la majeure partie de la représentation de connaissances conceptuellement orientée consiste dans le fait que, pour des raisons d'opportunité technique, la représentation de concepts en termes prototypiques et la possibilité d'exploiter des formes de raisonnement conceptuel fondé sur la typicalité n'est pas permis. En revanche, dans les sciences cognitives, certaines données en faveur de concepts prototypiques et de formes non-monotoniques de raisonnement conceptuel ont été étudiées de façon exhaustive. Ce gap cognitif-representationnel et de raisonnement représente un problème pour les systèmes computationnels, puisque l'information prototypique joue un rôle crucial dans plusieurs tâches importantes. Dans la lignée des theories "dual process" du raisonnement et de la rationalité, nous soutenons que la représentation conceptuelle dans les systèmes computationnels devrait être confiée à (au moins) deux composantes représentationnelles, chacune spécialisée dans le traitement de différents genres de processus de raisonnement. Dans cet article, nous présentons les avantages computationnels de cette approche "dual process", en les comparant avec d'autres solutions adoptées pour traiter de ce problème.

1 Introduction

In this article we concentrate on the problem of representing concepts in the context of artificial intelligence (AI) and of computational modeling of cognition. It is a problem that has a great relevance, for example, in the field of the development of computational ontologies.

One of the main problems of most contemporary concept oriented knowledge representation systems (including formal ontologies), is represented by the fact that they do not allow, for technical convenience, the representation of concepts in prototypical terms nor the possibility of exploiting forms of typicality-based conceptual reasoning. Conversely, in cognitive sciences, evidence exists in favor of prototypical concepts and non-monotonic forms of, approximate, conceptual reasoning have been extensively studied (see sect. 2 below). This "cognitive" representational and reasoning gap constitutes a problem for the computational systems, since prototypical information plays a crucial role in many relevant tasks. The historical reasons concerning the motivations of the abandon, in AI, of the typicality-based systems in favor of more rigorous formal approaches is sketched in section 3.

Given this state of affairs, we propose that some suggestions to face this problem should come from the psychology of reasoning. Indeed, in our view, a mature methodology to approach knowledge representation should take advantage also from the empirical results of cognitive research. In this paper we propose an approach to conceptual representation inspired by the so-called *dual process theories* of reasoning and rationality [Stanovich & West 2000], [Evan & Frankish 2008] (sect. 4). According to such theories, the existence of two different types of cognitive systems is assumed. The systems of the first type (type 1) are phylogenetically older, unconscious, automatic, associative, parallel and fast. The systems of the second type (type 2) are more recent, conscious, sequential and slow, and are based on explicit rule following. In our opinion, there are good *prima facie* reasons to believe that, in human subjects, tasks usually accounted by the above mentioned KR systems are tasks of the type 2 (they are difficult, slow, sequential tasks), while exceptions and prototypical knowledge play an important role in processes such as categorization, which is more likely to be a type 1 task: it is fast, automatic, and so on. Therefore, we advance the hypothesis that conceptual representation in computational systems should be equipped by (at least) two different kind of components[1] each responsible for different processes: type 2 processes, involved in complex inference tasks and which does not take into account the representation of prototypical knowledge, and fast and automatic type 1 processes, which perform categorization taking advantage from prototypical information associated to concepts (sect. 5).

2 Prototypical effects vs. compositionality in concept representation

In the field of cognitive psychology, most research on concepts moves from the critiques to the so-called classical theory of concepts, i.e. the traditional point of view according to which concepts can be defined in terms of necessary and sufficient conditions. The central claim of the classical theory of concepts is that every concept *c* is defined in terms of a set of features (or conditions) *f*1, ..., *fn* that are individually necessary and jointly sufficient for the application of *c*. In other words, everything that satisfies features *f*1, ..., *fn* is a *c*, and if anything is a *c*, then it must satisfy *f*1, ..., *fn*. For example, the features that define the concept *bachelor* could be *human*, *male*, *adult* and *not married*; the conditions defining *square* could be *regular polygon* and *quadrilateral*. This point of view was unanimously and tacitly accepted by psychologists, philosophers and linguists until the middle of the 20th century.

[1] For a similar approach, see [Piccinini 2011]. A way to split the traditional notion of concept along different lines has been proposed by [Machery 2005].

Chronologically, the first critique of classical theory was due to a philosopher: in a well known section from the *Philosophical Investigations*, Ludwig Wittgenstein observes that it is impossible to identify a set of necessary and sufficient conditions to define a concept such as GAME [Wittgenstein 1953, § 66]. Therefore, concepts exist which cannot be defined according to classical theory, i.e. in terms of necessary and sufficient conditions. Concepts such as GAME rest on a complex network of *family resemblances*. Wittgenstein introduces this notion in another passage in the *Investigations*: «I can think of no better expression to characterise these similarities than "family resemblances"; for the various resemblances between members of a family: build, features, colour of eyes, gait, temperament, etc. etc.» (*ibid*., § 67).

Wittgenstein's considerations were corroborated by empirical psychological research: starting from the seminal work by Eleanor Rosch [Rosch 1975], with the psychological experiments that showed how common-sense concepts do not obey the requirement of the classical theory[2]: common-sense concepts cannot usually be defined in terms of necessary and sufficient conditions (and even if for some concepts such a definition is available, subjects do not use it in many cognitive tasks). Concepts exhibit *prototypical effects*: some members of a category are considered better instances than others. For example, a robin is considered a better example of the category of birds than, say, a penguin or an ostrich. More central instances share certain typical features (e.g. the ability of flying for birds, having fur for mammals) that, in general, are neither necessary nor sufficient conditions.

Prototypical effects are a well established empirical phenomenon. However, the characterisation of concepts in prototypical terms is difficult to reconcile with the *compositionality* requirement. In a compositional system of representations, we can distinguish between a set of *primitive*, or *atomic, symbols* and a set of *complex symbols*. Complex symbols are generated from primitive symbols through the application of a set of suitable recursive syntactic rules (generally, a potentially infinite set of complex symbols can be generated from a finite set of primitive symbols). Natural languages are the paradigmatic example of compositional systems: primitive symbols correspond to the elements of the lexicon, and complex symbols include the (potentially infinite) set of all sentences.

In compositional systems, the meaning of a complex symbol *s* functionally depends on the syntactic structure of *s* as well as the meaning of primitive symbols in it. In other words, the meaning of complex symbols can be determined by means of recursive semantic rules that work in parallel with syntactic composition rules. This is the so-called *principle of compositionality of meaning*, which Gottlob Frege identified as one of the main features of human natural languages.

Within cognitive science, it is often assumed that concepts are the components of thought, and that mental representations are compositional structures recursively built up starting from (atomic) concepts. However, according to a well known argument by [Jerry Fodor 1981], prototypical effects cannot be accommodated with compositionality. In brief, Fodor's argument runs as follows: consider a concept like PET FISH. It results from the composition of the concept PET as well as the concept FISH. However, the prototype of PET FISH cannot result from the composition of the prototypes of PET and FISH. For example, a typical PET is furry and warm, a typical FISH is greyish, but a typical PET FISH is neither furry and warm nor greyish. Therefore, some strain exists between the requirement of compositionality and the need to characterise concepts in compositional terms.

3 Representing concepts in computational systems

The situation outlined in the section above is, to some extent, reflected by the state of the art in AI and, in general, in the field of computational modelling of cognition. This research area often seems to hesitate between different (and hardly compatible) points of view [Frixione & Lieto 2011]. In AI, the representation of concepts is faced mainly within the field of knowledge representation

[2] On the empirical inadequacy of the classical theory and the psychological theories of concepts see [Murphy 2002].

(KR). Symbolic KR systems (KRs) are formalisms whose structure is, broadly speaking, language-like. This usually entails assuming that KRs are compositional.

In their early development (historically corresponding to the late 1960s and the 1970s), many KRs oriented to conceptual representations attempted to take into account suggestions from psychological research. Examples are early semantic networks and frame systems. Frame and semantic networks were originally proposed as alternatives to the use of logic in KR. The notion of frame was developed by Marvin Minsky [1975] as a solution to the problem of representing structured knowledge in AI systems[3]. Both frames and most semantic networks allowed for the possibility to characterise concepts in terms of prototypical information.

However, such early KRs were usually characterised in a rather rough and imprecise way. They lacked a clear formal definition, with the study of their meta-theoretical properties being almost impossible. When AI practitioners tried to provide a stronger formal foundation to concept oriented KRs, it turned out to be difficult to reconcile compositionality and prototypical representations. As a consequence, they often chose to sacrifice the latter.

In particular, this is the solution adopted in a class of concept-oriented KRs which were (and still are) widespread within AI, namely the class of formalisms that stem from the so-called structured inheritance networks and the KL-ONE system [Brachman & Schmolze 1985]. Such systems were subsequently called terminological logics, and today are usually known as *description logics* (DLs) [Baader et al. 2010]. From a formal point of view, description logics are subsets of first order predicate logic that, if compared to full first order logic, are computationally more efficient.

In more recent years, representation systems in this tradition (such as the formal ontologies) have been directly formulated as logical formalisms (the above mentioned description logics, [Baader et al. 2010]), in which Tarskian, compositional semantics is directly associated to the syntax of the language. This has been achieved at the cost of not allowing exceptions to inheritance and, in so doing, we have forsaken the possibility to represent concepts in prototypical terms. From this point of view, such formalisms can be seen as a revival of the classical theory of concepts, in spite of its empirical inadequacy in dealing with most common-sense concepts. Nowadays, DLs are widely adopted within many application fields, in particular within that of the representation of ontologies. This state of affairs does constitute a problem for KR systems since prototypical effects in categorisation and, in general, in category representation are the greatest importance in representing concepts in both natural and artificial systems.

Several proposal have been advanced, in order to extend concept-oriented KRs, and DLs in particular, in such a way to represent non-classical concepts. Various fuzzy extensions of DLs [Bobillo & Straccia 2009] and ontology oriented formalisms have been proposed in order to represent vague information in semantic languages. However, from the standpoint of conceptual knowledge representation, it is well known [Osherson & Smith 1981] that approaches to prototypical effects based on fuzzy logic encounter difficulties with compositionality. (In short, Osherson and Smith show that the approaches to prototypical effects based on fuzzy logic are vulnerable to the problem of compositionality mentioned at the end of section 2 above.)

A different way to face the representation non classical concepts in DL systems are DL extensions based non some non-monotonic logic. For example, [Baader & Hollunder 1995] proposed an extension of the ALCF system based on Reiter's default logic[4]. The same authors, however, point out both the semantic and computational difficulties of this integration and, for this reason, propose a restricted semantics for open default theories, in which the default rules are only applied to

[3] Many of the original articles describing these early KRs can be found in [Brachman & Levesque 1985], a collection of classic papers of the field.

[4] The authors pointed out that "Reiter's default rule approach seems to fit well into the philosophy of terminological systems because most of them already provide their users with a form of 'monotonic' rules. These rules can be considered as special default rules where the justifications – which make the behaviour of default rules non-monotonic – are absent".

individuals explicitly represented in the knowledge base. [Bonatti et al. 2006] proposed an extension of DLs with circumscription. One of the reasons for applying circumscription is the possibility to express prototypical properties with exceptions, something which is done by introducing "abnormality" predicates whose extension is minimized. More recently, [Giordano et al. 2013] proposed an approach to defeasible inheritance based on the introduction in the *ALC* DL of a typicality operator **T**, which, in part, allows to reason about prototypical properties and inheritance with exceptions. However, we shall return later (sect. 5) on the plausibility of non-monotonic extensions of DL formalisms, as a way to face the problem of representing concepts in prototypical terms.

4 The dual process approach and its computational developments

In our opinion, a different approach to face the above described state of affairs should come from the so-called dual process theories. As anticipated in the introductory section, according to the *dual process theories* [Stanovich & West 2000], [Evan & Frankish 2008], two different types of cognitive systems exist, which are called respectively system(s) 1 and system(s) 2.

System 1 processes are automatic. They are phylogenetically older, and are shared between humans and other animal species. They are innate, and control instinctive behaviors; so, they do not depend on training or specific individual abilities, and generally are cognitively undemanding. They are associative, and operate in a parallel and fast way. Moreover, system 1 processes are not consciously accessible to the subject.

System 2 processes are philogenetically more recent, and are specific to the human species. They are conscious and cognitively penetrable (i.e., accessible to consciousness), and are based on explicit rule following. As a consequence, if compared to system 1, system 2 processes are sequential and slower, and cognitively more demanding. Performances that depend on system 2 processes are usually affected by acquired skills and differences in individual capabilities.

The dual process approach was originally proposed to account for systematic errors in reasoning tasks: systematic reasoning errors (consider the classical examples of the selection task or the so-called conjunction fallacy) should be ascribed to fast, associative and automatic system 1 processes, while system 2 is responsible for the slow and cognitively demanding activity of producing answers that are correct with respect to the canons of normative rationality. An example is the well known Linda problem, in which participants are given a description of "Linda" that stresses her independence and liberal views, and then asked whether it is more likely that she is (a) a bank teller or (b) a bank teller and active in the feminist movement. Participants tend to choose (b), since it fits the description of Linda (following the 'representativeness heuristic'), even though the co-occurrence of two events cannot be more likely than one of them alone.

A first theoretical attempt to apply the dual process theory in the field of computational modelling has been developed by Sloman [Sloman 1996], whose proposal is based on the computational distinction between two types of reasoning systems. System 1 is associative and is attuned to encoding and processing statistical regularities and correlations in the environment. System 2 is rule based. The representations in this system are symbolic and unbounded, in that they are based on propositions that can be compositionally combined. Sloman uses Smolensky's [Smolensky 1988] connectionist framework to describe the computational differences between system 1 and system 2. Smolensky contrasted two types of inferential mechanisms within a connectionist framework: an intuitive processor and a conscious rule interpreter. Sloman claims that both system 1 (intuitive processor) and system 2 (conscious rule interpreter) are implemented by the same hardware but use different types of knowledge that are differently represented. The relationship between the systems is described as interactive. Moreover, he proposes that the two systems operate in concert and produce different outputs both useful in different ways. Therefore, by using the terminology proposed by Evans [Evans, 2008], in Sloman the two computational systems are supposed to be "parallel-competitive" in nature, differently from the traditional “default-interventionist” approach, which is typical of the dual process proposals (according to such “default-interventionist” approach the deliberative S2 reasoning processes can inhibit the biased responses of the S1 systems and replace them with “correct outputs” based on reflective reasoning).

In recent years, the cognitive modeling community posed a growing attention on the dual process theories as a framework for modeling cognition “beyond the rational”, in the sense of [Kennedy et al. 2012]. This determined two main effects: (i) a strong effort of rethinking some classical cognitive architectures in terms of the dual process theory; and (ii) the development of new cognitively inspired artificial systems embedding some theoretical aspects of the dual theory. In this section we will give review some examples of these two classes.

As far as point (i) is concerned, there are at least three examples of pre-existing hybrid cognitive architectures that have been reconsidered in terms of the dual process hypothesis, namely: Soar has recently included the initial System 1 form of assessment of a situation and used it as the basis for reinforcement learning [J.E. Laird 2008], ACT-R [Anderson et al. 2004] now integrates explicit, declarative (i.e. system 2) representations and implicit procedural (system 1) cognitive processes[5] and similarly the CLARION architecture [Hélie & Sun 2010] adopts a dual representation of knowledge, consisting of a symbolic component to manage explicit knowledge (system 2) and a low-level component to manage tacit knowledge (system 1). More recently, in the field of AGI (Artificial General Intelligence, see [Mc Carthy, 2007]) a dual process multi-purpose cognitive architecture has also been proposed [Strannegård et al. 2013]. The architecture is based on two memory systems: (i) a long-term memory, which is an autonomous system that develops automatically through interactions with the environment, and (ii) a working memory, which used to perform (resource-bounded) computation. Computations are defined as processes in which working memory content is transformed according to rules that are stored in the long-term memory. In such architecture, the long-term memory is modeled as a transparent neural network that develops autonomously by interacting with the environment and that is able to activate both S2 and S1 processes. The working memory (S1) is modeled as a buffer containing nodes of the long-term memory.

In the meanwhile of the mentioned developments within the field of cognitive architectures, some new systems have been recently proposed, which are directly inspired by the dual process approach. A first example is the mReasoner model [Khemlani & Johnson-Laird 2012], developed with the aim of providing a unified computational architecture of reasoning[6] based on the mental model theory proposed by Philip Johnson-Laird. The mReasoner architecture is based on three components: a system 0, a system 1 and a system 2. The last two correspond to those hypothesized by the dual process approach. System 0 operates at the level of linguistic pre-processing. It parses the premises of an argument using natural language processing techniques, and then creates an initial intensional model of them. System 1 uses this intensional representation to build an extensional model, and uses heuristics to provide rapid reasoning conclusions; finally, system 2 carries out more demanding processes to searches for alternative models, if the initial conclusion does not hold or if it is not satisfactory.

A second system has been proposed by [Laure et al. 2012]. The authors adopt an extended version of the dual process approach, based on the hypothesis that the system 2 is subdivided in two further levels, respectively called “algorithmic” and “reflective”. The goal of Laure and colleagues is to build a multi agent and multilevel architecture able to represent the emergence of emotions in a biologically inspired computational environment.

[5] Differently from CLARION, ACT-R does not use a double level, e.g. symbolic and sub-symbolic, of representations. Its “type 1” processes are based, as the “type 2” ones, on the same layer of procedural-based, symbolic, knowledge.

[6] The appeal to the need of unitary computational architectures in Cognitive Science and AI is not new. See e.g. Newell (1990) .

Another system which can be included into this class has been proposed by [Pilato et al. 2012]. The authors do not explicitly mention the dual process approach; however, they built a hybrid system for conversational agents (chatbots) where the agents' background knowledge is represented using both a symbolic and a subsymbolic approach. The authors associate different types of representations to different types of reasoning. Namely deterministic reasoning is associated to symbolic (system 2) representations, and associative reasoning is associated to the subsymbolic (system 1) component. Differently from the other systems following the dual approach, the authors do not make any claim about the sequence of activation and the conciliation strategy of the two representational and reasoning processes. Such conciliation strategy, on the other hand, plays a crucial role in the field of the dual-process based computational systems. Elsewhere [Frixione & Lieto 2012, 2014] we have presented a novel computational strategy for the integration of the S1 and S2 processes in the field of a dual process account of concepts in semantic technologies. Such strategy, differently from both the "default-interventionist" proposal (where S1 processes are the default ones and are then checked against the S2) and from the Sloman's proposal of "naturally-parallel" computations, is computationally more conservative and safe, since the typicality based reasoning is considered as an extension of the classical one and is exploited only in the case of unsatisfactory results provided by the classical, S2, component (that is compositional and that performs only deductive, and therefore logically correct, inferences).

It is worth noting that other examples could be found, of computational models that are in some sense akin to the dual process proposal, even if their proponents do not explicitly mention this approach. Consider for example many hybrid, symbolic-connectionist systems, in which the connectionist component is used to model fast, associative processes, while the symbolic component is responsible for explicit, declarative computations [see Frasconi et al. 2002].

5 Dual processes and concept representation

In our opinion, the distinction between system 1 and system 2 processes could be plausibly applied to the problem of conceptual representation as it emerged in the sections above. In particular, categorization based on prototypical information is in most cases a fast and automatic process, which does not require any explicit effort, and which therefore can presumably be attributed to a type 1 system. On the contrary, the types of inference that are typical of DL systems (such as classification and consistency checking) are slow, cognitively demanding processes that are more likely to be attributed to a type 2 system.

Let us consider for example the case of classification. In a DL system, classifying a concept in a taxonomy amounts to individuate its more specific superconcepts and its more general subconcepts. As an example, let us suppose that a certain concept C is described as a subconcept of the concept S, and that each instance of C has at least three fillers of the attribute R that are instances of the concept B. Let us assume also that these traits are conjunctly sufficient to be a C (i.e., everything that is an S with at least three fillers of the attribute R that are Bs is also a C). Let us suppose now that another concept C' is described as an S with exactly five fillers of the attribute R that are B's, and that B' is a subconcept of B. On the basis of these definitions, it follows that every C' must in its turn be also a C; in other terms, C' must be a subconcept of C. Classifying a concept amounts to identify such implicit superconcept-subconcept relations in a taxonomy. But for human subjects such a process is far from resulting natural, fast and automatic.

So, the inferential task of classifying concepts in taxonomies is *prima facie* qualitatively different from the task of categorizing items as instances of a certain class on the basis of typical traits (e.g., the task of categorizing Fido as a dog because he barks, has fur and wag his tail).

Note that, in this perspective, the approach to prototypical representation of concepts based on non-monotonic extensions of some DL formalism (see sect. 3 above) seem to be particularly implausible. The idea at the basis of such an approach is that prototypical representation of concepts should be obtained by augmenting DLs with non-monotonic constructs that should allow to represent defeasible information. In such a way, categorization based on prototypical traits results to be a process homogeneous to classification, but still more demanding, and carried out with a still more complex formalism (it is well known that, in general, non-monotonic formalism have worst computational properties if compared to their monotonic counterparts)[7].

In this spirit, we argue that conceptual representation in computational systems could be demanded to (at least) two different kind of components responsible for different processes: type 2 processes, involved in complex inference tasks and which does not take into account the representation of prototypical knowledge, and fast and automatic type 1 processes, which perform such tasks as categorization taking advantage from prototypical information associated to concepts. Moreover, it is likely that, in human mind, prototypical information about concepts is coded in different ways [Murphy 2002, Machery, 2005].

Recently, an implementation of the presented dual process-conceptual proposal has been realized [Ghignone, Lieto and Radicioni 2013] and preliminarily tested in a knowledge-based system involved in a question-answering task. In such system, imprecise and common sense natural language descriptions of a given concept were provided as queries. The task designed for the evaluation consisted in individuating the appropriate concept that fits a given description by exploiting the inferential capability of the proposed hybrid conceptual architecture. According to the assumption presented in [Frixione and Lieto 2013], the S1 component is based on the Conceptual Spaces framework [Gärdenfors 2000] and the classical S2 component on standard Description Logics and ontology based formalisms. An example of such common-sense descriptions is: “the big carnivore with black and yellow stripes" denoting the concept of tiger. The obtained preliminary results are encouraging and show that the identification and retrieval of concepts described with typical features is considerably improved by such hybrid architecture w.r.t. the classical case, based simply on the use of ontological knowledge. Furthermore, this result is obtained with a relatively limited computational effort if compared with the other, logic-based, approaches. This results suggest that a dual process approach to conceptual representation of concepts can be beneficial for enhancing the performance of artificial systems in tasks involving non-classical conceptual reasoning.

References

Baader, Franz, Calvanese, Diego, McGuinness, Deborah, Nardi, Daniele and Patel-Schneider, Peter, *The Description Logic Handbook: Theory, Implementations and Applications* (2nd edition); Cambridge: Cambridge University Press, 2010 .

Bonatti, Piero, Lutz, Carsten and Wolter, Franz, Description logics with circumscription, *Proc. KR 2006*, 400–410, 2006.

Brachman, Ronald and Levesque, Hector 1985 *Readings in Knowledge Representation*, Los Altos, CA: Morgan Kaufmann.

Brachman, Ronald and Schmolze, James G., An overview of the KL-ONE knowledge representation system, *Cognitive Science*, 9, 171-216, 1985.

[7] This does not amount to claim that, in general, non-monotonic extensions of DLs are useless. Our claim is simply that they seem to be unsuitable (and cognitively implausible) for the task of representing concepts in prototypical terms.

Evans, Jonathan and Frankish, Keith (eds.) *In Two Minds: Dual Processes and Beyond*. New York, NY: Oxford UP, 2008.

Evans, Jonathan, 'Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition. *Annual Review of Psychology* 59 (2008): 255–78, 2008.

Frasconi, Paolo, Gori, Marco, Kurfess, Franz J., Sperduti, Alessandro, Editorial: Special issue on integration of symbolic and connectionist systems, in "*Cognitive Systems Research*", 3, 121-123, 2002.

Frixione, Marcello, Lieto, Antonio, Representing Concepts in Artificial Systems: A Clash of Requirements, *Proceedings of the 4th International HCP11 Workshop*, 75-82, 2011.

Frixione, Marcello, Lieto, Antonio, *Representing Concepts in Formal Ontologies: Compositionality vs Typicality Effects*, in Logic and Logical Philosophy, vol. 21 No. 4, December 2012, 391-414, 2012.

Frixione, Marcello, Lieto, Antonio, Dealing with Concepts: from Cognitive Psychology to Knowledge Representation, *Frontiers in Psychological and Behavioural Science*, Vol. 2, Iss. 3, Jul. 2013, 96-106, 2013.

Frixione, Marcello, Lieto, Antonio, Towards an Extended Model of Conceptual Representations in Formal Ontologies: A Typicality-Based Proposal, in *Journal of Universal Computer Science*, Vol. 20 (3), 2014.

Fodor, Jerry, The present status of the innateness controversy, J. Fodor, *Representations*, Cambridge, MA: MIT Press, 1981.

Gärdenfors, Peter, *Conceptual Spaces: The Geometry of Thought*. The MIT Press/Bradford Books. 2000.

Ghignone, Leo, Lieto, Antonio, Radicioni, Daniele, Typicality-based Inference by Plugging Conceptual Spaces into Ontologies, in *Proceedings of AIC 2013, International Workshop on Artificial Intelligence and Cognition*, 68-79, 2013.

Giordano, Laura, Gliozzi, Valentina, Pozzato, Gianluca, Olivetti, Nicola, A Non-monotonic Description Logic for reasoning about typicality, *Artificial Intelligence* 165-202, 2013.

Hélie, Sébastien and Sun, Ron, Incubation, insight, and creative problem solving: a unified theory and a connectionist model, *Psychological Review*, 117(3), 994–1024, 2010.

Kennedy, William G., Ritter, Frank E., Lebiere, Christian, Juvina, Ion, Oltramari, Alessandro, Grateh, Jonathan and Young, Richard M., *ICCM Symposium on Cognitive Modeling of Processes "Beyond Rational*", Proceedings of ICCM 2012 11th International Conference on Cognitive Modeling. 2012.

Khemlani, Sangeet and Johnson-Laird, Philip N., The processes of inference, *Argument and Computation*, 0(0): 1-17, 2012.

Laird, John E., Extending the Soar Cognitive Architecture, *In Proceedings of the Conference on Artificial General Intelligence,* ISO Press, Memphis, 2008.

Larue, Othalia, Poirier, Pierre and Nkambou, Roger, Emotional emergence in a symbolic dynamical architecture, *Proc. BICA 2012,* Palermo, Italy, 2012.

Machery, Edouard, Concepts are not a Natural Kinds, *Philosophy of Science*. 2005.

Mc Carthy, John, From here to human-level AI, Artificial Intelligence, 171, 1174-1182. 2007.

Minsky, Marvin, A framework for representing knowledge. *The Psychology of Computer Vision*, P. H. Winston (ed.), McGraw-Hill. 1975.

Murphy, Gregory, *The Big Book of Concepts*, Cambridge, MA: The MIT Press. 2002.

Newell, Allen, *Unified Theory of Cognition*, Cambridge, MA: Harvard University Press. 1990.

Piccinini, Gualtiero, Two Kinds of Concept: Implicit and Explicit, *Dialogue*, 50, 179-193. 2011.

Pilato, Giovanni, Augello, Agnese and Gaglio, Salvatore, A modular system oriented to the design of versatile knowledge bases for chatbots, *ISRN Artificial Intelligence*, vol. 2012, Article ID 363840, 10 pages, doi:10.5402/2012/363840, 2012.

Rosh, Eleanor, Cognitive representation of semantic categories, *Journal of Experimental Psychology* 104: 573-605. 1975.

Sloman, Steven A., The Empirical Case for Two Systems of Reasoning, *Psychological Bulletin* 119 (1996),3-22, 1996.

Smolensky, Paul, On the proper treatment of connectionism. *Behavioral and Brain Sciences, 11, 1-23*, 1988.

Stanovich, Keith E. and West, Richard, Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences* 23, 5: 645- 65. 2000.

Strannegård, Claes, von Haugwitz, Rickard, Wessberg, Johan, Balkenius Christian, A Cognitive Architecture Based on Dual Process Theory, *Artificial General Intelligence, Lecture Notes in Computer Science* Volume 7999, 140-149, 2013.

Wittgenstein, Ludwig, *Philosophische Untersuchungen*. Oxford: Blackwell. 1953.