Philosophia Scientiæ

Travaux d'histoire et de philosophie des sciences

18-3 | 2014

Logic and Philosophy of Science in Nancy (I)

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

Marcello Frixione and Antonio Lieto

Electronic version

URL: <http://journals.openedition.org/philosophiascientiae/1005>

DOI: 10.4000/philosophiascientiae.1005

ISSN: 1775-4283

Publisher

Éditions Kimé

Printed version

Date of publication: 1 October 2014

Number of pages: 139-152

ISBN: 978-2-84174-689-7

ISSN: 1281-2463

Electronic reference

Marcello Frixione and Antonio Lieto, « Formal Ontologies and Semantic Technologies: A “Dual Process” Proposal for Concept Representation », *Philosophia Scientiæ* [Online], 18-3 | 2014, Online since 19 January 2015, connection on 02 November 2020. URL : <http://journals.openedition.org/philosophiascientiae/1005> ; DOI : <https://doi.org/10.4000/philosophiascientiae.1005>

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

Marcello Frixione

DAFIST, University of Genova (Italy)

Antonio Lieto

University of Torino – ICAR-CNR, Palermo (Italy)

Résumé : Pour la plupart des systèmes de représentation de la connaissance orientés concept, l'un des problèmes principaux relève de la commodité technique. À savoir, la représentation de connaissance en termes prototypiques, tout comme la possibilité d'exploiter des formes de raisonnement conceptuel basées sur la typicalité, ne sont pas autorisées. Au contraire, dans les sciences cognitives, il existe des données en faveur de concepts prototypiques, et des formes non-monotoniques de raisonnement conceptuel ont été largement étudiées. Ce fossé cognitif concernant la représentation et le raisonnement constitue un problème pour les systèmes computationnels, puisque l'information prototypique joue un rôle crucial dans plusieurs tâches importantes. Dans la lignée des théories du raisonnement et de la rationalité dites du double processus, nous soutenons que la représentation conceptuelle dans les systèmes computationnels devrait dépendre de (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 en termes de double processus, et les comparons brièvement avec d'autres solutions d'orientation logique, adoptées pour traiter du même problème.

Abstract: One of the main problems of most contemporary concept-oriented knowledge representation systems is one of technical convenience. Namely the representation of knowledge in prototypical terms and the possibility of exploiting forms of typicality-based conceptual reasoning, are not permitted. In contrast, in the cognitive sciences, evidence exists in favour of prototypical

concepts, and non-monotonic forms of conceptual reasoning have been extensively studied. This “cognitive” representational and reasoning gap constitutes a problem for computational systems, since prototypical information plays a crucial role in many relevant tasks. Inspired by the so-called dual process theories of reasoning and rationality, we propose that conceptual representation in computational systems should rely on (at least) two representational components, that are specialized in dealing with different kinds of reasoning processes. In this article, the theoretical and computational advantages of such “dual process” proposals are presented and briefly compared to other logic-oriented solutions adopted to confront the same problem.

1 Introduction

In this article we concentrate on the problem of representing concepts in the context of artificial intelligence (AI) and of computational modelling of cognition. These are highly relevant problems, for example in the field of the development of computational ontologies.

One of the main problems of most contemporary concept oriented knowledge representation systems (KRs, including formal ontologies), is that, for technical convenience, they do not admit the representation of concepts in prototypical terms. In this way, the possibility of exploiting forms of typicality-based conceptual reasoning is excluded. In contrast, in the cognitive sciences evidence exists in favour of prototypical concepts, and non-monotonic forms of approximate conceptual reasoning have been extensively studied (see section 2, below). This “cognitive” representational and reasoning gap constitutes a problem for 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 favour of more rigorous formal approaches is outlined 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 (KR) should also take advantage of the empirical results of cognitive research. In this paper, we put forward an approach towards conceptual representation inspired by the so-called *dual process theories* of reasoning and rationality [Stanovich & West 2000], [Evans & Frankish 2008] (section 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 than those of the second type, 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 for by KRs are type 2 tasks (they are difficult, slow, sequential tasks). However, exceptions and prototypical knowledge could play an important role in processes such as categorisation, 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 kinds of component¹ each responsible for different processes: type 2 processes, involved in complex inference tasks and that do not take into account the representation of prototypical knowledge, and fast automatic type 1 processes, which perform categorisation taking advantage of prototypical information associated with concepts (section 5).

2 Prototypical effects vs. compositionality in concept representation

In the field of cognitive psychology, most research on concepts moves from 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 can be defined in terms of a set of features (or conditions) $f_1, \dots, f_n$ that are individually necessary and jointly sufficient for the application of c . In other words, everything that satisfies features $f_1, \dots, f_n$ is a c , and if anything is a c , then it must satisfy $f_1, \dots, f_n$. For example, the features that define the concept *bachelor* could be *human*, *male*, *adult* and *unmarried*; 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.

Chronologically, the first critique of classical theory was by 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. [Wittgenstein 1953, §67]

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].

Wittgenstein's considerations were corroborated by empirical psychological research. Starting from the seminal work by Eleanor Rosch [Rosch 1975], psychological experiments 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 to fly for birds, to have fur for mammals) that, in general, are neither necessary nor sufficient conditions for the concept.

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 is 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 on the meaning of primitive symbols within *s*. 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 [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 and the concept FISH. However, the prototype of PET FISH cannot result from the composition of the prototypes of PET and FISH. Indeed, 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.

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

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 oscillate between different (and hardly compatible) points of view [Frixione & Lieto 2011]. In AI, the representation of concepts lies mainly within the KR field. Symbolic 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 that are oriented towards 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 [Minsky 1975] as a solution to the problem of representing structured knowledge in AI systems.³ Both frames and most semantic networks allowed the possibility of concept characterisation 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 for concept-oriented KRs, it turned out to be difficult to reconcile compositionality and prototypical representations. In consequence, practitioners 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, Calvanese *et al.* 2010]. From a formal point of view, DLs 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 DLs, [Baader, Calvanese *et al.* 2010]), in which Tarskian, compositional semantics is directly associated with the syntax of the language. This has been achieved at the cost of not allowing exceptions to inheritance and, in this way, we have forsaken the possibility to represent concepts in prototypical terms. From this point of view, such formalisms can be seen

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.

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 fields of application, in particular within that of the representation of ontologies. This is a problem for KRs since prototypical effects in categorisation and, in general, in category representation, are of the greatest importance in representing concepts in both natural and artificial systems.

Several proposals have been advanced to extend concept-oriented KRs and DLs in particular, in such a way as to represent non-classical concepts. Various fuzzy extensions of DLs [Bobillo & Straccia 2009] and ontology-oriented formalisms have been proposed 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.

A different way to face the representation of non-classical concepts in DL systems exists, namely DL extensions based on 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 restricted semantics for open default theories, in which the default rules are only applied to individuals explicitly represented in the knowledge base. [Bonatti, Lutz *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, Gliozzi *et al.* 2013] proposed an approach to unfeasible inheritance based on the introduction in the *ALC* DL of a typicality operator *T*, which allows prototypical properties and inheritance with exceptions to be reasoned about in part. However, we shall return later (section 5) to the plausibility of non-monotonic extensions of DL formalisms, as a way to confront the problem of representing concepts in prototypical terms.

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".

4 The dual process approach and its computational developments

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

System 1 processes are automatic. They are phylogenetically the older of the two, and are often 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 are generally cognitively undemanding. They are associative, and operate in a parallel and fast way. Moreover, they are not consciously accessible to the subject.

System 2 processes are phylogenetically more recent than system 1 processes, 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, 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 “heuristic representativeness”), 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 to 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 sys-

tem 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 represented differently. The relationship between the systems is described as interactive. Moreover, he proposes that the two systems operate in concert and produce different outputs that are both useful but 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, that is typical of the dual process proposals (according to such “default-interventionist” approach the deliberative system 2 reasoning processes can inhibit the biased responses of the system 1 processes and replace them with “correct outputs” based on reflective reasoning).

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

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. Soar has recently included the initial system 1 form of assessment of a situation and used it as the basis for reinforcement learning [Laird 2008], ACT-R [Anderson, Bothell *et al.* 2004] now integrates explicit, declarative (i.e., system 2) representations and implicit procedural (system 1) cognitive processes.⁵ 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 [McCarthy 2007]) a dual process multi-purpose cognitive architecture has been proposed [Strannegård, von Haugwitz *et al.* 2013]. The architecture is based on two memory systems: (i) long-term memory, which is an autonomous system that develops automatically through interactions with the environment, and (ii) working memory, which is 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 modelled as a transparent neural network that develops autonomously

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.

by interacting with the environment and that is able to activate both system 1 and system 2 processes. The working memory (system 1) is modelled as a buffer containing nodes of the long-term memory.

At the same time as the above-mentioned developments within the field of cognitive architectures, some new models were proposed, that are directly inspired by the dual process approach. A first example is the mReasoner model [Khemlani & Johnson-Laird 2013], developed with the aim of providing a unified computational architecture of reasoning⁶ 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 search for alternative models if the initial conclusion does not hold or if it is not satisfactory.

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

Another model that can be included in this class has been proposed by [Pilato, Augello *et al.* 2012]. These 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 with different types of reasoning. Namely deterministic reasoning is associated with symbolic (system 2) representations, and associative reasoning is linked to the subsymbolic (system 1) component. Differently from the other models that follow 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. However, such a conciliation strategy 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 system 1 and system 2 processes in the field of a dual process account of concepts in semantic technologies. Such a strategy, differently from both the “default-interventionist” proposal (where system 1 processes are the default ones and are then checked against the system 2) and from Sloman’s proposal of “naturally-parallel” com-

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

putations, is computationally more conservative and safe, since the typicality based reasoning is considered as an extension of the classical one and is only exploited 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 of computational models that are in some sense akin to the dual process proposal can be found 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 [Wermter & Sun 2000].

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, categorisation 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. In contrast, 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 taxonomy amounts to individualising 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 in conjunction are sufficient to be a C (i.e., everything that is an S with at least three fillers of the attribute R that are B s 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 identifying such implicit superconcept-subconcept relations in taxonomy. But for human subjects such a process is far from a natural, fast and automatic one.

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

Note that in this perspective, the approach to prototypical representation of concepts based on non-monotonic extensions of some DL formalism (see section 3 above) seems to be particularly implausible. The idea at the basis of

such an approach is that the prototypical representation of concepts should be obtained by increasing DLs with non-monotonic constructs that should allow defeasible information to be represented. In such a way, the categorisation based on prototypical traits is a process homogeneous to classification, but still more demanding, and needs to be carried out with an even more complex formalism (it is well-known that, in general, non-monotonic formalisms have worse computational properties than their monotonic counterparts).⁷

In this spirit, we argue that conceptual representation in computational systems could demand (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 automatic type 1 processes, which perform such tasks as categorisation taking advantage of prototypical information associated with concepts. Moreover, it is likely that, in the human mind prototypical information about concepts is coded in different ways [Murphy 2002; Machery 2005].

Recently, an implementation of the dual process-conceptual proposal presented was achieved [Ghignone, Lieto *et al.* 2013] and preliminary tests were carried out in a knowledge-based system involved in a question-answering task. In such a system, imprecise and common sense natural language descriptions of a given concept were provided as queries. The task designed for the evaluation consisted of individualising 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 & Lieto 2013], the system 1 component is based on the Conceptual Spaces framework [Gärdenfors 2000] and the classical system 2 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 preliminary results obtained are encouraging and show that the identification and retrieval of concepts described by typical features is considerably improved using such hybrid architecture, with respect to the classical case, based simply on the use of ontological knowledge. Furthermore, this result is obtained with a relatively limited computational effort compared to the other, logic-based, approaches. These results suggest that a dual process approach to conceptual representation of concepts can be beneficial to enhance the performance of artificial systems in tasks involving non-classical conceptual reasoning.

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.

Bibliography

- ANDERSON, John R., BOTHELL, Daniel, BYRNE, Michael D., DOUGLASS, Scott, LEBIERE, Christian, & QIN, Yulin [2004], An integrated theory of the mind, *Psychological Review*, 111(4), 1036–1060, doi:10.1037/0033-295X.111.4.1036.
- BAADER, Franz, CALVANESE, Diego, MCGUINNESS, Deborah, NARDI, Daniele, & PATEL-SCHNEIDER, Peter [2010], *The Description Logic Handbook: Theory, Implementations and Applications*, Cambridge: Cambridge University Press, 2nd edn.
- BAADER, Franz & HOLLUNDER, Bernhard [1995], Embedding defaults into terminological knowledge representation formalisms, *Journal of Automated Reasoning*, 14(1), 149–180, doi:10.1007/BF00883932.
- BOBILLO, Fernando & STRACCIA, Umberto [2009], An owl ontology for fuzzy owl 2, in: *Foundations of Intelligent Systems – 18th International Symposium (ISMIS 2009)*, edited by J. Rauch, Z.W. Raś, P. Berka, & T. Elomaa, Berlin Heidelberg: Springer, Lecture Notes in Computer Science, vol. 5722, 151–160, doi:10.1007/978-3-642-04125-9_18.
- BONATTI, Piero, LUTZ, Carsten, & WOLTER, Franz [2006], Description logics with circumscription, in: *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning*, 400–410.
- BRACHMAN, Ronald & LEVESQUE, Hector [1985], *Readings in Knowledge Representation*, Los Altos: Morgan Kaufmann.
- BRACHMAN, Ronald & SCHMOLZE, James G. [1985], An overview of the KL-ONE knowledge representation system, *Cognitive Science*, 9, 171–216, doi:10.1016/S0364-0213(85)80014-8.
- EVANS, Jonathan [2008], Dual-processing accounts of reasoning, judgment, and social cognition, *Annual Review of Psychology*, 59, 255–278, doi:10.1146/annurev.psych.59.103006.093629.
- EVANS, Jonathan & FRANKISH, Keith (eds.) [2008], *In Two Minds: Dual Processes and Beyond*, New York: Oxford University Press.
- FODOR, Jerry [1981], The present status of the innateness controversy, in: *Representations*, edited by J. Fodor, Cambridge, MA: MIT Press, 257–316.
- FRIXIONE, Marcello & LIETO, Antonio [2011], Representing concepts in artificial systems: A clash of requirements, in: *Proceedings of the 4th International HCP11 Workshop*, 75–82.

- [2012], Representing concepts in formal ontologies. compositionality vs. typicality effects, *Logic and Logical Philosophy*, 21(4), 391–414, doi:10.12775/LLP.2012.018.
- [2013], Dealing with concepts: From cognitive psychology to knowledge representation, *Frontiers in Psychological and Behavioural Science*, 2(3), 96–106.
- [2014], Towards an extended model of conceptual representations in formal ontologies: A typicality-based proposal, *Journal of Universal Computer Science*, 20(3), 257–276.
- GÄRDENFORS, Peter [2000], *Conceptual Spaces: The Geometry of Thought*, Cambridge, MA: MIT Press.
- GHIGNONE, Leo, LIETO, Antonio, & RADICIONI, Daniele [2013], Typicality-based inference by plugging conceptual spaces into ontologies, in: *Proceedings of AIC 2013, International Workshop on Artificial Intelligence and Cognition*, 68–79.
- GIORDANO, Laura, GLIOZZI, Valentina, POZZATO, Gianluca, & OLIVETTI, Nicola [2013], A non-monotonic description logic for reasoning about typicality, *Artificial Intelligence*, 195(0), 165–202, doi:10.1016/j.artint.2012.10.004.
- HÉLIE, Sébastien & SUN, Ron [2010], Incubation, insight, and creative problem solving: A unified theory and a connectionist model, *Psychological Review*, 117(3), 994–1024, doi:10.1037/a0019532.
- KENNEDY, William G. *et al.* [2012], Symposium: Cognitive modeling of processes “Beyond Rational”, in: *Proceedings of ICCM 2012 11th International Conference on Cognitive Modeling*, Berlin: Universitätsverlag der TU Berlin, 55–58.
- KHEMLANI, Sangeet & JOHNSON-LAIRD, Philip N. [2013], The processes of inference, *Argument & Computation*, 4(1), 4–20, doi:10.1080/19462166.2012.674060.
- LAIRD, John E. [2008], Extending the soar cognitive architecture, in: *Proceedings of the Conference on Artificial General Intelligence*, Amsterdam: IOS Press, 224–235.
- LARUE, Othalia, POIRIER, Pierre, & NKAMBOU, Roger [2012], Emotional emergence in a symbolic dynamical architecture, in: *BICA'12*, Palermo, 199–204.
- MACHERY, Edouard [2005], Concepts are not a natural kinds, *Philosophy of Science*, 72, 444–467.

- McCARTHY, John [2007], From here to human-level AI, *Artificial Intelligence*, 171(18), 1174–1182.
- MINSKY, Marvin [1975], A framework for representing knowledge, in: *The Psychology of Computer Vision*, edited by P. H. Winston, New York: McGraw-Hill, 211–277.
- MURPHY, Gregory [2002], *The Big Book of Concepts*, Cambridge, MA: MIT Press.
- NEWELL, Allen [1990], *Unified Theory of Cognition*, Cambridge, MA: Harvard University Press.
- OSHERSON, Daniel N. & SMITH, Edward E. [1981], On the adequacy of prototype theory as a theory of concepts, *Cognition*, 9(1), 35–58, doi: 10.1016/0010-0277(81)90013-5.
- PICCININI, Gualtiero [2011], Two kinds of concept: Implicit and explicit, *Dialogue*, 50(1), 179–193, doi:10.1017/S0012217311000187.
- PILATO, Giovanni, AUGELLO, Agnese, & GAGLIO, Salvatore [2012], A modular system oriented to the design of versatile knowledge bases for chatbots, doi: 10.5402/2012/363840.
- ROSCH, Eleanor [1975], Cognitive representation of semantic categories, *Journal of Experimental Psychology*, 104(3), 573–605.
- SLOMAN, Steven A. [1996], The empirical case for two systems of reasoning, *Psychological Bulletin*, 119, 3–22.
- SMOLENSKY, Paul [1988], On the proper treatment of connectionism, *Behavioral and Brain Sciences*, 11, 1–23.
- STANOVICH, Keith E. & WEST, Richard [2000], Individual differences in reasoning: Implications for the rationality debate?, *Behavioral and Brain Sciences*, 23(5), 645–665.
- STRANNEGÅRD, Claes, VON HAUGWITZ, Rickard, WESSBERG, Johan, & BALKENIUS, Christian [2013], A cognitive architecture based on dual process theory, in: *Artificial General Intelligence*, edited by Kai-Uwe Kühnberger, Sebastian Rudolph, & Pei Wang, Berlin; Heidelberg: Springer, Lecture Notes in Computer Science, vol. 7999, 140–149, doi:10.1007/978-3-642-39521-5_15.
- WERMTER, Stefan & SUN, Ron [2000], *Hybrid Neural Systems*, Heidelberg; New York: Springer.
- WITTGENSTEIN, Ludwig [1953], *Philosophische Untersuchungen*, Oxford: Blackwell.