How Psychology and Cognition can inform the Creation of Ontologies in Semantic Technologies.

Paula C. ENGELBRECHT[1&2], Itiel E. DROR[2]
[1]*Ordnance Survey, Romsey Road, Southampton, SO16 4GU, UK*
[2]*School of Psychology, University of Southampton, Southampton, SO17 1BJ, UK*

This paper demonstrates how cognitive psychology can contribute to the development of ontologies for semantic technologies and the semantic web. Specifically, the way in which the human cognitive system structures and processes conceptual information can act as a model for structuring formal ontologies, and can guide knowledge elicitation and the use of controlled natural languages. This paper concludes that cognitive psychology and the science of ontology development can inform each other: One practical implication is that ontology developers need to be more specific and explicit about what they mean by the term concept when explaining the use of concepts in ontologies to domain experts. The science of ontology development in turn can suggest new avenues for psychological enquiry.

Introduction

In computer science, ontologies are formal descriptions of sets of concepts related to specific domains [1]. They play a key role in semantic technologies and the semantic web; and are typically expressed using formal logic. The knowledge representation language for the Semantic Web, for example, is based on a family of logics (called description logics) that are decidable fragments of first-order logic [2]. In creating ontologies, ontology developers are faced with a variety of theoretical and practical problems. These include the representational formalisms for encoding ontologies, how to scope them and how they should be scaled and segmented [3].

This paper demonstrates how cognitive psychology can play a critical role in dealing with some of these issues. Specifically, it focuses on how our understanding of human knowledge representation and information processing could contribute to three aspects of ontology development. The first aspect to be discussed is the knowledge elicitation stage of ontology development during which domain experts act as knowledge sources. It is argued that knowledge engineers need to take into consideration both the strengths and the limitations of the human cognitive system in order to generate domain ontologies that are fit for purpose. The second aspect is concerned with the use of controlled natural languages to express ontologies. Evidence is presented about how the pragmatics of day-to-day language use (specifically, about how people use and understand generic and universal statements) can affect the content of ontologies authored by domain experts. The third and final aspect illustrates how human knowledge representation can inform the structuring of ontologies.

1. Capturing knowledge representations (the human as an informant)

Domain ontologies are external representations of experts' subject-related knowledge. Knowledge engineers extract this domain knowledge using a variety of established knowledge elicitation techniques such as interviews, concept mapping and card sorting [4, 5]. Due to its basis in formal logic, there is an inherent assumption within the discipline of ontology development that conceptual knowledge comprises relatively fixed and stable mental representations which can be accessed using the right knowledge elicitation techniques. This view is reflected in traditional approaches in cognitive psychology which assume that concepts are static mental representations that are retrieved from long-term memory when required [6]. However, there is evidence to suggest that internal representations of concepts are not stable entities but are the product of dynamic, context dependent processes [7]. This claim is supported by experimental findings which illustrate that people's conceptualisations of categories vary as a function of context and are unstable across time [7, 8, 9]. For example, it has been demonstrated that depending on the task context, category construction can be based on taxonomic (e.g. a sparrow is a bird) or thematic (e.g. cake and candles belong to the birthday party schema) relationships between concepts. In other words, depending on the context a "dog" can be classified as a mammal, a pet, a guide to the blind, a guard or a friend. This classification in turn will affect which "dog" attributes will come to mind. When thinking about dogs in the context of "kinds of mammals" the fact that dogs can serve as guides to the blind is less likely to come to mind than the fact that they bear life young.

The flexibility of internal representations of conceptual knowledge has several implications for knowledge elicitation. First, knowledge elicitation, rather than being a means of retrieving fixed concepts from memory, is an interactive process during which domain ontologies are modified and constructed (at least in part). This means that the construction of a domain ontology is of benefit to the domain expert, because engaging in this activity can solidify vague ideas and lead to novel insights. Second, the fact that internal knowledge representations are not fixed highlights the importance of setting the right context during knowledge elicitation. Concepts can encompass both context sensitive and context independent information [7]. Context-independent information about concepts (e.g. that robins are red-breasted) is considered to be highly accessible and relatively stable whereas context-dependent information is less accessible [7]. Knowledge elicitation exercises that are carried out in a neutral context are likely to elicit context-independent information. This might be desirable if one wishes to generate ontologies that can be used and reused across a wide variety of tasks and contexts. Frequently however, (especially in the case of semantic technologies) ontologies are intended to be used for specific tasks. On these occasions it is important to render the task context salient (for example by the way in which interview questions are worded) in order for the concepts and concept attributes that are most essential to the task to be elicited.

Third, in addition to contextual factors, other cognitive processes also affect knowledge representations. For example, a classical finding in cognitive psychology shows that humans can only hold seven plus or minus two chunks of information in working memory at any given point in time [10]. Working memory capacity can be extended quite easily, however, by means of external aids. An implication of this for knowledge elicitation is that it is advisable to have written or visual representations of the discussed information available to both the knowledge engineer and the domain expert. The cognitive processes that come into play during knowledge elicitation can be very subtle. For example, it has been found that the process of comparing two similar categories (e.g. a donkey and a horse) leads to an increase in the perceived similarity between them even when differences are listed [11]. This finding would suggest that certain knowledge elicitation methods, such as card sorting (which involves the comparison of concepts in order to group them), may lead to the omission of defining attributes that are not shared with the comparison category. In order to capture domain ontologies that are fit for purpose knowledge engineers need to consider the above issues during knowledge elicitation.

2. Knowledge representation (the domain expert as knowledge engineer)

Structured natural languages such as Sydney OWL Syntax [12], Attempto Controlled English [13] and Rabbit [14] are subsets of natural languages with restricted grammars and vocabularies. They enable domain experts to author ontologies without having to deal with the complexities of formal logic that underlie computational ontologies. Structured natural languages are relatively easy to learn because they harness people's pre-existing knowledge in language comprehension and production. An unintended consequence of this is that people are also likely to apply the same processes and assumptions that they use in day-to-day language use. A mismatch between how people understand structured natural languages and the assumptions inherent in formal logic can give rise to computational ontologies that contain erroneous statements; i.e. statements that do not comply with formal logic. This point is illustrated by considering the use of the universal quantifier in structured natural languages, and how "and" in English must sometimes be interpreted as "or" in formal logic.

Universal statements are generic assertions (e.g. "ducks lay eggs") to which a universal quantifier (e.g. "all", "every") is added. Structured natural languages use universal quantifiers to denote inheritance (e.g. "Every Man is a kind of Human") and other relationships that apply to all members of a category. Generics statements (e.g. "tigers are striped") express generalisations but lack quantifiers (e.g. all, some or most). A recent study on how people interpret generic and universal assertions [15] distinguishes three different types of statements: the first type consists of *characteristic* statements (that describe typical concept attributes), e.g. "ducks lay eggs", the second type refers to *striking* consequences, e.g. "ticks carry Lyme disease" and the third to statements that refer to the *majority* of instances, e.g. "cars have radios". The study described in [15] found that people tend to agree with generic statements that refer to characteristic properties (agreement on 89% of occasions) and, to a lesser extent, with generic statements that describe striking and majority properties (agreement on 68% of occasions). For universal statements, hardly any agreement was found for striking and majority properties (agreement on only 7% of occasions). However, participants had a tendency to agree with universal statements that describe characteristic properties, e.g., "all lions have manes" (agreement on 47% of occasions), even though these statements were in fact false.

According to Khemlani and his colleagues [15], understanding generic characteristic statements is a cognitively primitive operation which is less demanding on the cognitive system than dealing with quantified assertions. The authors argue that people tend to process universal characteristic statements like generic characteristic statements. This argument, combined with the observation that counter examples to characteristic statements (e.g. “female lions do not have manes”) are not readily accessible, helps to explain the observed tendency to agree with universal characteristic statements.

The study described above deals with language comprehension rather than production. Nevertheless, although the empirical question remains to be addressed, there is no reason to assume that the same processes would not apply to an ontology authoring process where people generate the characteristic examples themselves. The findings thus suggest that the use of the universal quantifier in structured natural languages, although it avoids a certain proportion of errors (namely those caused by the erroneous belief that striking and majority properties refer to all category instances), does not circumvent people’s tendency to agree with universal statements about characteristic properties. This tendency might cause domain experts to develop ontologies that contain erroneous concepts. That is concepts which contain universal statements that do not apply to all instances of that concept (“all ducks lay eggs” for example).

There are ways of minimizing this potential source of “erroneous” concepts. For example, an ontology authoring tool which supports ontology authoring using structured natural languages could have inbuilt probe questions, e.g., “Can you think of a duck which does not lay eggs?”, to prompt the domain expert to access counter examples. Alternatively, the above findings could give rise to a reconsideration of the nature of ontologies. What can be said with certainty about all instances of a given category, and how meaningful and useful is this information? With the possible exception of inheritance, exceptions to almost every statement which initially seems true can be found. It might therefore be useful to consider how modal logic (which can express necessity and possibility) or default logic (which can express facts that are true in the majority of cases) could be combined with the description logics currently used in semantic web ontology languages like OWL. An early example of this is [16]. The utility of this would partly depend on the kind of use to which an ontology is put.

An awareness of how people deal with universal statements and other related issues should inform the development of languages and tools designed to develop domain ontologies. Another area of cognitive psychology that plays a central role in this context is research on errors and biases in human reasoning. For example, it has been found that people’s performance on syllogistic reasoning problems is affected by the believability of the conclusion [17]. This finding illustrates that humans find it hard to disregard their existing knowledge when performing abstract reasoning tasks. The manner in which domain experts represent and access conceptual knowledge and reason about it affects the ontologies they develop. It follows that what is known about natural language use and human reasoning should inform the development of both controlled natural languages and the ontology authoring tools in which they are implemented.

3. Logical ontologies (the domain expert as a model)

The most basic relationship in ontologies is the taxonomic (IS-A) relationship which denotes class inheritance, and has an analogue with property inheritance. Ontologies often contain taxonomic hierarchies of concepts [1]. Organising concepts into class hierarchies is both elegant and economic because relationships that apply to several related categories do not have to be repeated over and over again. For example, if one knows that all mammals breathe one does not need to encode this information again for cats; it can be inferred from the class hierarchy. Similarly, within the cognitive psychology literature it has been argued that conceptual knowledge is represented in the form of inheritance hierarchies in long-term memory [18, 19]. According to this view the hierarchical structure represented in memory consists of IS-A links [20]. The model predicts that it should take people longer to infer that ‘a salmon is an animal’ than ‘a salmon is a fish’. This is because the verification of the former statement supposedly requires the traversal of several IS-A links, whereas the latter requires traversal of only one. Evidence in support of this hypothesis has been found for both category inclusion decisions and for the verification of property statements, e.g., a salmon can swim versus a salmon has gills [19].

An alternative interpretation of the above findings is that they reflect a computational process. This view assumes that because closely related concepts share many properties, the relationships between them are easier to compute [20]. The finding that it takes people longer to verify statements containing atypical than typical category members; e.g. it takes them longer to verify that a penguin is a type of bird than that a robin is a type of bird, lends support to this interpretation [20]. These and other findings have given rise to the conclusion that IS-A relationships arise from a set of computations rather than being a central component of semantic long-term memory [20, 21]. The argument that human conceptual knowledge is not organised taxonomically is further strengthened by evidence (discussed in a previous section) which shows that concepts are context-sensitive representations that are assembled dynamically when needed [7, 22].

The above discussion indicates that to model conceptual ontologies more closely to human conceptual representations requires shallow rather than deep hierarchies. The term shallow hierarchy refers to concept hierarchies with few taxonomic levels. Shallow hierarchies are less supportive of subclasses inheriting the properties of their superclasses than deep hierarchies simply because there are fewer taxonomic levels from which properties can be inherited. This means that the relevant attributes need to be associated to concepts directly rather than being inherited from concepts further up in the hierarchy. The use of shallow taxonomic hierarchies has several advantages in ontology construction. For example, although humans can happily categorise an instance as belonging to more than one category (e.g. a dog can be both an animal and a friend) best practice in ontology design does not encourage the explicit specification of multiple inheritance relationships; however, it is possible to give a category both animal-like and friend-like attributes. A further argument in favour of shallow hierarchies is that ontologies – as they are currently implemented – do not allow for non-monotonic inheritance. When making "dog" a subclass of "friend", unwanted inherited properties of the "friend" category cannot be overridden. In sum, the psychological literature makes some suggestions about how to structure domain ontologies.

4. Discussion and Conclusions

This paper has highlighted two ways in which what is known about cognition and cognitive psychology can inform ontology development. First, several activities in ontology development (e.g. knowledge elicitation and ontology authoring) involve domain experts. The efficiency with which these activities are carried out and the utility of the resulting ontologies can be improved by considering human information processing and its limitations. Second, the human cognitive system in general, and human knowledge representation in particular, can act as a model for the structure of ontologies.

Cognitive psychology can make suggestions for improvements in how ontologies are developed, structured and used. For example, the above discussion illustrates that ontology developers need to be more explicit and specific about what a concept is. One possible solution would be to define a concept as something that holds true for all normal category members. This would avoid the interchangeable use of *characteristic* (typical) concepts and *majority* (true in most cases) concepts in ontologies. Doing so can help avoid non-trivial problems further down the line when ontologies are used for automatic reasoning; this is especially crucial when ontologies which make different assumptions are merged.

Another problem of ontology development that has been highlighted in the above discussion is scope. How does one decide which concepts and concept attributes to include in an ontology and which to leave out? Psychological theories of conceptual coherence can bear on this issue. It has been argued that conceptual coherence arises from peoples' theories about the world [23]. For example, people are much more likely to list "does not fly" as an attribute of "penguins" than "sharks" because peoples knowledge of birds would predict that penguins (but not sharks) can fly. Thus peoples theories about the world determine which attributes are important to a concept and which are not. As discussed above, peoples concepts of a given category are unstable, context dependent, constructs. The fact that different theories about the world are salient at different times helps to explain this. An implication for ontology development is that the scope and purpose section of a domain ontology should constrain its theoretical framework. By doing so, the decision of which concept and concept attributes to include (and which to leave out) is made much easier.

A related problem to the above is how does one decide at which level of granularity an ontology should be captured? Research in cognitive psychology has shown that humans favour basic level categories when thinking and talking about the world. This basic level of abstraction is a "compromise between the accuracy of classification at a maximally general level and the predictive power of a maximally specific level" [20]. In other words, basic level concepts are those which optimize both informativeness and distinctiveness [20]. Take for example the concept "collie" (which is subordinate to the concept "dog"). Although "collie" is more informative than "dog", it is also less distinctive (it is harder to tell a "Collie" from a "German Sheperd" than it is to tell a "dog" from a "cat"). The concept "animal" (which is a super-ordinate of "dog"), on the other hand, is highly distinctive (it is easy to tell animals apart from other concepts at the same level such as "plants" and "objects") but not very informative. Based on these observations the most appropriate recommendation might be to represent ontologies at the basic level and to move only one level up or down from it. Doing so will ensure that the hierarchy of the ontology is kept shallow. What constitutes the basic level of abstraction for a given ontology is very much dependent on the purpose for the ontology. For example, within the biology domain the concept "amino acid" would be basic to an ontology describing proteins but not in an ontology that captures "cellular processes".

Ontologies are intended for use both by computers and humans, yet they represent knowledge differently. For example, human conceptual representations are flexible and dynamic whereas logical ontologies are relatively fixed constructs. Furthermore, whereas human knowledge representations can deal well with uncertain and ambiguous class memberships logical ontologies (as they are currently implemented) cannot. Due to these and other inherent differences in the way in which formal logic represents knowledge and the way in which humans do, ontology development requires a careful balance between the needs of the human and the needs of the machine. In order to find optimal solutions to these problems, ontology developers need to be knowledgeable of the limitations of both humans and of formal logic.

References

[1] Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5, 199-220.

[2] Smith, M. K., Welty, C. & McGuinness, D. L. (2004b), Owl web ontology language guide, W3C Recommendation 20040210, W3C. Latest version available at http://www.w3.org/TR/owlguide/.

[3] Seidenberg, J. & Rector, A. (2006). Web ontology Segmentation: analysis, classification and use. In *Proceedings of the 15th International Conference on World Wide Web*, 18-22. New York: ACM Press.

[4] Garcia Castro, A., Rocca-Serra, P., Stevens, R., Taylor, C., Nashar, K., Ragan, M. A. & Sansone, S. A. (2006). The use of concept maps during knowledge elicitation in ontology development processes – the nutrigenomics use case. *BMC Bioinformatics*, *7*, 267.

[5] Wang Y., Sure, Y., Stevens, R. & Rector, A. (2006). Knowledge elicitation plug-in for protégé: card sorting and laddering. In R. Mizoguchi, Z. Shi & F. Giuchiglia, *First Asian Semantic Web Conference (ASWC'06), 4185* of LNCS, pp. 552-565.

[6] Rummelhart, D. E. & McClelland, J. D. (1986). *Explorations in the microstructure of congnition, volume 1: Foundations.* Cambridge MA: Bradford.

[7] Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization.* New York: Cambridge University Press.

[8] Roth, E. M. & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, 15, 346-378.

[9] Lin, E. L. & Murphy, G. L. (2001). Thematic relations in adults' concepts. *Journal of Experimental Psychology: General*, 130, 3-28.

[10] Miller, G. A. (1956). The magical number seven plus or minus two: some limits on our capacity for processing information. *The Psychological Review*, 63, 81-97.

[11] Boroditsky, L. (2007). Comparison and the development of knowledge. *Cognition*, *102*, 118-128.

[12] Cregan, A., Schwitter, R. & Meyer, T. (2007). Sydney OWL Syntax - towards a Controlled Natural Language Syntax for OWL 1.1. In *Proceedings of OWLED 2007 Workshop on OWL: Experiences and Directions*, vol, 258 of CEUR Workshop Proceedings. http://ceur-ws.org/

[13] Kaljurand, K. & Fuchs, N. E. (2007). Verbalizing OWL in Attempto Controlled English. In *Proceedings of OWLED 2007 Workshop on OWL: Experiences and Directions*, vol, 258 of CEUR Workshop Proceedings. http://ceur-ws.org/

[14] Hart, G., Dolbear, C. & Goodwin, J. (2007). Lege Feliciter: Using structured English to represent a topographic hydrology ontology. In *Proceedings of OWLED 2007 Workshop on OWL: Experiences and Directions*, vol, 258 of CEUR Workshop Proceedings. http://ceur-ws.org/

[15] Khemlani, S., Glucksberg, S. & Rubio Fernandez, P. (2007). Do ducks lay eggs? How people interpret generic assertions. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society*, 64-70. Austin, TX: Cognitive Science Society

[16] Klinov, P & Parsia, B. (2008). Demonstrating Pronto: a Non-Monotonic Probabilistic OWL Reasoner. OWL Experiences and Directions DC 2008, Washington DC 1-2 April 2008. http://www.webont.org/owled/2008dc/papers/owled2008dc_paper_2.pdf

[17] Evans, J. St. B. T., Barston, J. L. & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition*, 11, 295-306.

[18] Anderson, J. R. & Bower, G. H. (1973). *Human associative memory*. Washington DC: Winston.

[19] Collins, A. M. & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Behavior and Verbal Learning*, 23, 625-642.

[20] Murphy, G. L. (2002). *The big book of concepts*. Massachusetts: MIT Press.

[21] Sloman, S. A. (1998). Categorical inference is not a tree: the myth of inheritance hierarchies. *Cognitive Psychology*, 35, 1-33.

[22] Smolensky, P. (1991). Connectionism, constituency and the language of thought. In B. Loewer & G. Rey (Eds.), *Meaning in Mind: Fodor and his Critics.* Oxford: Basil Blackwell.

[23] Murphy, G. L. & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.