Ontological Properties of Animals in a Children's Dictionary With and Without Common-Sense Knowledge

Julia M. Taylor (jtaylor@purdue.edu)
Computer and Information Technology & CERIAS, Purdue University
401 N. Grant Street, W. Lafayette, IN 47907-2021 USA

Victor Raskin (vraskin@purdue.edu)
Linguistics & CERIAS, Purdue University
656 Oval Drive, W. Lafayette, IN 47907-2086 USA

Christian F. Hempelmann (chempelm@purdue.edu)
Linguistics and CERIAS, Purdue University
656 Oval Drive, W. Lafayette, IN 47907-2086 USA

Abstract

The paper applies a limited version of the resources of Ontological Semantic Technology to the descriptions of animals in the American Heritage First Dictionary and constructs a partial ontology from them. The explicitly mentioned properties in the descriptions are then supplemented by common-sense knowledge that the descriptions assume available to their young readership, and the output is compared to the previous one. The results, albeit modest, shed some interesting light on the most similar and dissimilar pairs of animals, as described in text.

Keywords: common-sense knowledge; Ontological Semantic Technology; children's dictionary; animal dataset; similarity.

Introduction

The paper explores the common-sense knowledge that is necessary to fully understand the definitions/descriptions (henceforth, just descriptions) of around 100 animals in the 2007 edition of the American Heritage First Dictionary (AHFD 2007) aimed at children in grades K-2 (ages 5-8). It does not address common-sense reasoning. The purpose is to get a grasp on how implicit information affects the structure of perceived knowledge, in this case of animal descriptions, and similarity among entities.

The AHFD contains about 2,000 entries, claimed, almost entirely correctly, to be written with a controlled vocabulary so that that every description contains only words that also have entries in the dictionary. What is challenging in this design for our task is the possible implication of self-sufficiency, that is, of a much reduced dependency on the child's knowledge of the world, unstated explicitly in the natural language descriptions but (unconsciously) assumed to be present for full comprehension.

Most AFHD definitions for animals follow the "genus proximum, differentia specifica" format that is common for dictionaries, in particular for the classification of animal species: "Goldfish is a kind of *fish* [genus]. Goldfish are usually *small* and *orange* [differentia]." Apart from the differentiating specifics (small, orange) and the few properties that are specifics for the genus (in the AHFD, fish live in water, have tails, can swim well) and are thus inherited, the remaining knowledge necessary to understand such definitions remains implicit, because it is presumed to be common-sense of different kinds. There is, for example, no mention of a fish's gills or fins.

Capturing common-sense knowledge is a daunting task. We are assuming that descriptions of the (animal) world of this dictionary requires less common-sense knowledge simply because this world is more restricted than that of an adult and the dictionary was obviously designed to accommodate that. If that is so, then getting a grasp of that knowledge may be more feasible than in case of a common, unlimited, adult-level natural language communication[1].

The goal is, then, to illustrate how much information is lost when common-sense knowledge is not made explicit. Using the methods of computational semantics, specifically our Ontological Semantic Technology, we are taking advantage of the unique design of the dictionary to identify the required common-sense knowledge for a reasonably full comprehension of its animal descriptions. In this way, we aim to get a sense of its common-sense knowledge dependency. As a result, we also hope to clarify some issues concerning the very nature of common-sense knowledge and the feasibility of its computational acquisition and use, which is, as a matter of act, our primary and real concern.

In Section 1, we introduce the notions of 'hard' and 'soft' common-sense knowledge and explore its relation to underdetermination of reality by language and to saliency and, then, to ontology and natural language meaning, contingency, and instantiation.. In Section 2, we will briefly survey pertinent prior work. Section 3 will sketch out the Ontological Semantic Technology, our research tool as we applied it to the material. Section 4 compares the worldview on the animals that the descriptions define with the one complemented by the common-sense knowledge necessary to understand them. Section 5 discusses the results,

[1] The distinction between children as "novices" who know less about many domains than "expert" adults is well established (e.g., Carey 1985)-for better or worse.

identifies the strengths and weaknesses of our approach, and discusses the future lines of research.

Kinds of Common-Sense Knowledge

Hard and Soft Common-Sense Knowledge

We are introducing this new pair of terms to differentiate between two kinds of common-sense knowledge that a reader of the AHFD must possess to fully comprehend a description. If that reader does not understand a word in the description and that word has its own AHFD entry, he or she may access the required knowledge from that entry, where it is explicitly stated. So, from the point of the initial entry, this information is implied but from the point of the dictionary, it is explicitly stated. We call this information the 'soft common-sense knowledge.' If, on the contrary, some information that is needed for a full comprehension of the entry is not stated explicitly anywhere in the dictionary, we refer to it as the 'hard common-sense knowledge.' The paper focuses on the latter.

Starting with the randomly selected AHFD description of *snake*, we line up (see graph in http://web.ics.purdue.edu /~vraskin/snake_new_label.pdf) the lexical chains underlying the entry dependency: if entry E(x) uses word y in the description of word x, and y is not previously mentioned in the chain, E(x) leads to E(y). If E(i) does not evoke any new words in its description, it becomes a terminal in the dependency line. The longest, 10-node dependency line holds, starting with the topmost leftmost node and ending with the rightmost node at the bottom of the picture: SNAKE is-a REPTILE is-a ANIMAL is-a-not PLANT agent-of MAKE has-agent BEE agent-of FILL result-in FULL precondition-of-not HOLD unspecified ROOM. What this perfectly representative branch illustrates is that there is no consistent or predictable semantic dependency in the chain and that the vagaries of lexicographic connection can traverse the domain of knowledge, common-sense and other, in all directions, with some connections not easily explained.

Altogether, the knowledge required to understand every word in the description of *snake* as well as every word in the descriptions of those words, and, in turn, every word in the descriptions of those words, and so on to the end of the chain, is expressed in 86 entries. Realistically speaking, no 5-year-old will read all the entries: much more likely, they will have the requisite knowledge of the words. Nevertheless, this information is made available by the AHFD compilers, perhaps similarly to the glosses, footnotes, and explanatory appendices in adult-level materials. Its availability makes it not quite common-sense knowledge—so we refer to this explicit, but remote information, as weak common-sense knowledge.

Underdetermination and Saliency

It is known that language underdetermines reality (see, for instance, Barwise and Perry 1983; Nirenburg and Raskin 2004): no matter how fine-grained or verbose the description of an event, there will be tons of details about the situation that will remain unmentioned. If two men walk into the room, a report of that may include what they look like, what they wear, the speed of their movement, etc. But it will mention nothing about their places of birth, parents' names and occupations, what cars they drive, what they had for breakfast, etc.

Now, all that knowledge exists, and common-sense knowledge includes that these people have a birth place, have parents, likely drive cars (especially if they're Americans), etc. What is essential, however, is that most of the existing but implicit information is not prominent: much more likely, the prominence goes with the purpose of those people's entrance into the room, whether there is any cause for alarm or displeasure, etc. The amount of prominent, or salient common- sense knowledge is much more limited in any situation.

Unfortunately, saliency (see Giora 2003: 13-38 and references there) is dynamic and fluctuates very rapidly. In AHFD, however, saliency may be conveniently seen as deliberately delimited by the availability of entries for words, thus reflecting the compilers' notion of the mental model for a five-year-old's world.

Instantiation and contingency

In Ontological Semantic Technology (OST), the ontology consists of concepts and relations between them that are determined by properties. The concepts anchor lexical senses that are defined in the separate lexicon. Thus, one sense of the word *cat* is anchored in the ontological concept CAT. In a sentence, *A cat can jump from the floor to the top of a bookcase*, CAT is what the word *cat* means, i.e., a generic, any member of the class.

In the sentence, *Kisa the cat can jump from the floor to the top of a six-foot tall bookcase*, however, it is no longer a generic cat, but a specific instance of the concept, and the relationship between the meaning of the word *cat* and the ontological concept CAT is no longer that of generic anchoring. This instantiation makes the sentence contingent on a number of indices, such as the identities of the speaker and hearer, time, place, etc. (see Lewis 1972—cf. Bar Hillel's 1954 comment on rare non-contingent sentences, such as, *Ice floats on water*).

We understand common-sense knowledge as non-contingent and involving concepts, not their instances. It is about what exists in the world, not what we know about particular objects or events. Our common-sense knowledge includes the fact that houses may be painted in various colors; it does not include the fact that Tom's house is grey with burgundy trim.

So the common-sense knowledge left implicit by the AHFD is strong, non-contingent, and definitely less salient than the knowledge explicitly supplied by the AHFD in its descriptions.

Prior Pertinent Work on Common-Sense Knowledge

Distinguishing common-sense knowledge from other implicit types of knowledge has been an issue in approaches to knowledge engineering, and while it always is a central one, it often remains implicit. Knowledge-based NLP has (re-)matured enough both to be able to need as well as to accommodate the type of "deep" knowledge that overlaps with the varying notions of common sense.

McCarthy (1959) is often cited as the earliest mention of common sense in the literature, but Bar-Hillel's (1954) well-known example, "Little John played in his pen," is already a clear indication of the necessity and importance of the common-sense knowledge—in this case, about relative sizes of objects.

Prominently, Lenat (1990) started an early large-scale systematic project on acquisition of common-sense knowledge, CyC. His method was hand-coding by a large number of research engineers, with a high turnaround and no well-defined acquisition methodology, which affected results and rendered them unusable for the NLP community.

Gordon and Schubert's overview (2010) classifies current approaches to common-sense knowledge acquisition as: hand-authoring of rules, as in CyC; abstracting from clusters of propositions (e.g., Van Durme 2009); and directly interpreting general statements, such as glosses in dictionaries (e.g., Clark et al. 2008), akin to the approach of the present paper. Other researchers have used tagging, annotating, and/or generic machine learning techniques for automatically extracting implied common-sense knowledge from explicit text on the Web, about which Lin et al. (2004) have legitimate reservations, because explicit statements on the Web do not necessarily express common-sense knowledge.

Finally, we need to mention the area of research on common sense dedicated to children's development of such knowledge, not least related to their overall linguistic-cognitive development. In particular, children's knowledge about animals is one of the applications. Results that inform our present approach include that children focus on external features rather than internal organs, on habitats, on behavior relevant for humans (dangerous, edible) rather than cladistic accuracy ("Is a camel an ungulate?"), and that children's knowledge is derived from observations as much as instructions, parents, or media (see Prokop et al. (2007), Tunnicliffe et al. (2007), Byrne et al. (2010)).

In our own previous work (Taylor et al. 2011a), we include in the common-sense knowledge rules of a separate resource the knowledge-of-the-world information that is not already contained in the ontology and lexicon (see next section). in the experiment there, we processed text with our system, and as part of routine quality assurance, added the necessary common-sense knowledge wherever we failed to interpret the text correctly because of the unavailability of this information in our resources (after we have excluded other, more banal reasons for the failure, such as an error in the resources or a bug in the software). Thus, we identified as missing, for example, size classes necessary to understand spatial relations between physical objects, such as the understanding that a containing object should have greater dimensions than the (solid) object it contains.

In contrast to previous work, which addressed the identification and acquisition of common-sense knowledge by OST for the general purpose of processing text, this paper applies an appropriately limited version of our resources to a very limited corpus of a specific genre in an attempt to compare the ontological information following from the AHFD descriptions only with the ontological information arising from the descriptions supplemented by the common-sense knowledge that the descriptions imply in their readership.

Brief Introduction to OST

Charniak's (1972) often (mis)cited children's story is used primarily to discuss inferencing and, hence, reasoning. It is even more suitable for exemplifying (in square brackets) the most common common-sense knowledge that OST has to deal with in order to fulfill its function of representing the meaning of natural language text accurately and comprehensively.

Jane was invited to Jack's birthday party. [One brings presents to a birthday party. <u>Presents</u> are often purchased. To purchase something, one needs money.] *She wondered if he would like a kite. She went into her room and shook her piggy bank.* [Piggy banks contains <u>money</u>, usually coins. <u>Coins</u> make <u>noise</u> when shaken] *It made no sound.* [Coins make noise.] →(either there was no money in the piggy bank or just no coins but rather bills→ in the former case, Jane may have lacked the money to buy the present)].

The italicized part is the original story; our formulation of the common-sense knowledge is in square brackets; the parenthesized part following the first arrow represents our formulation of inferences in reasoning, and while definitely pertinent to common-sense knowledge, it will be left out of this paper. It is noteworthy that the reasoning statements are contingent on the story while common-sense knowledge is generic.

The first and essential function of OST is to interpret the text of the story. The OST processor reads each sentence linearly and looks it up, word by word, in the OST English lexicon. Every sense of every (non-auxiliary, non-parametric) word in the lexicon is anchored in an ontological concept, with its properties and fillers, and the fillers can be restricted by the sense. The OST ontology, unlike its lexicons, is language-independent (see Nirenburg and Raskin 2004 for the basic theory of Ontological Semantics, and Raskin et al. 2010, Hempelmann et al. 2010, Taylor and Raskin 2011, Taylor et al. 2010, 2011a,b, for the much revised OST).

To use a greatly simplified example, the sense of the English word *invite* will be anchored in the ontological concept, probably also labeled "INVITE." The label does not contain any but distinguishing information for the computer and can be any ASCII combination—it is there just for the convenience of the human acquirer.

INVITE

is-a	communicative event
agent	human
beneficiary	human
theme	social-gathering
purpose	entertainment

invite

Invite

agent	[preceding NP]
beneficiary	[following NP]
theme	[to NP]
…	

And the text meaning representation (TMR) of the first sentence of the story will result from matching the meaning of the NPs in the appropriate EVENT slots. The reality is, of course, harder, with more complex syntax, ambiguity, etc. The unenhanced-OST problem with the story is still more advanced: while TMR for each sentence is not hard to produce, the system will not be able to relate the sentences to each other, and the text will lack cohesiveness.

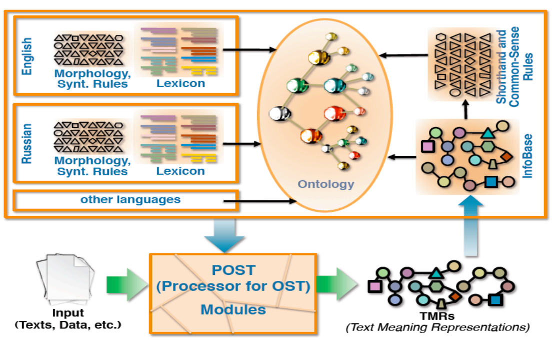

Figure 1: OST Architecture

In OST, the information processed prior to computing the TMR of the current sentence is used to clarify, complement, and disambiguate the current representation process. In this case, that information would be helpless and useless because, other than Jane as the agent, no previous sentence in the story even mentions objects in the following sentences, and it is the common objects (or events) that the anaphora/coreference resolution establishes as bridges between and among sentences. Jane will emerge from the story, as interpreted by the unenhanced OST, as performing three unrelated actions. It is the common-sense knowledge statements in the square brackets above that have to provide such common objects to make OST processing possible: the bridge words are underlined in the story above., and the common-sense knowledge enhanced text can be processed by OST normally.

This is why we recently added to the OST architecture (Figure 1 above) the common-sense knowledge resource (Taylor et al. 2011a) and the methodology of adding to it when the TMRs fall short of the (often hypothetical) gold standard (cf. Allen et al. 2008). .

Ontology of Descriptions and Ontology With Common-Sense Knowledge

In general, we are interested not only in reading and understanding a text, but also in structuring information that this text contains, as well as enhancing our ontology when newly acquired information requires. We are using information about animals from AHFD to see whether such task is possible. We then check whether supplying additional information (common-sense knowledge left implicit in the dictionary) would help with the task (cf. Perfors et al. 2005; Kemp et al. 2006).

Typically, a hierarchy is perceived as one of the most important properties in ontology construction. All animal descriptions of the dictionary provide such information. Unfortunately, sometimes a word is used that may have multiple senses (such as *cat* being a domestic cat or feline) thus creating a flawed hierarchy. One of the goals, then, is to identify such descriptions.

The proposed measure is conditional on an accepted membership assumption. If we assume the veracity of "B is A" as a reference point, which gives us a certain amount of knowledge about B in terms of its properties, we estimate the extent to which "C is B" is (dis)confirmed as $\frac{\sum_i 2^{-n} * w * hier_{P_i(C,B)}}{num(i)}$ where w is a property weight, and

$$hier_{P_i(C,B)} = \begin{cases} 1, & \mathrm{P.b} \in \mathrm{D_B}\ \&\ \mathrm{P.c} \in \mathrm{D_C}\ \&\ \mathrm{b} = \mathrm{c} \\ -1, & \mathrm{P.b} \in \mathrm{D_B}\ \&\ \mathrm{P.c} \in \mathrm{D_C}\ \&\ \mathrm{b} \neq \mathrm{c} \\ 0.1, & \mathrm{P} \in \mathrm{D_C}\ \&\ \mathrm{P} \notin \mathrm{D_B} \\ 0, & \text{otherwise} \end{cases}$$

Placement in the hierarchy as well as concepts' properties (may) affect similarity between concepts. For the purposes of this paper, we assumed that the properties that are taken into account are all equally weighted. We measure similarity of two concepts as $\frac{\sum_i 2^{-n} * w * sim_{P_i(A,B)}}{num(i)}$ where $sim_{Pi(A,B)}$ is defined as:

$$sim_{P_i(A,B)} = \begin{cases} 1, & \mathrm{P.a} \in \mathrm{D_A}\ \&\ \mathrm{P.b} \in \mathrm{D_B}\ \&\ \mathrm{a} = \mathrm{b} \\ -1, & \mathrm{P.a} \in \mathrm{D_A}\ \&\ \mathrm{P.b} \in \mathrm{D_B}\ \&\ \mathrm{a} \neq \mathrm{b} \\ 0.1, & \mathrm{P} \in \mathrm{D_A}\ \&\ \mathrm{P} \notin \mathrm{D_B} \parallel \mathrm{P} \notin \mathrm{D_A}\ \&\ \mathrm{P} \in \mathrm{D_B} \\ 0, & \text{otherwise} \end{cases}$$

Results and Conclusion

We first wanted to see what kind of structure we would get from the descriptions without the use of common sense. We calculated pair-wise similarity measurements for all animals with AHFD descriptions. The similarities ranged from -1.25 to 0.78. It is possible for the similarity to be –N where N is the number of properties in both descriptions and all properties in the descriptions match but their fillers do not. Having calculated the mean and standard deviation, we looked at the results that were at least 3 standard deviations away from the mean as most similar cases and most dissimilar ones. The dissimilar pairs were: ant/chicken, ant/crocodile, ant/pony, ant/whale, bee/chicken, beetle/chicken, bug/chicken, bug/shark, bug/whale, butterfly/chicken, caterpillar/chicken, caterpillar/crow, caterpillar/whale, chicken/cricket, chicken/fly, chicken/mosquito, chicken/moth, chicken/whale, cricket/whale, crocodile/whale, mosquito/whale, moth/shark, moth/whale, turtle/whale.

It should be noted that, with the exception of the chicken/whale, turtle/whale, and crocodile/whale pairs, the dissimilar pairs contain insects. One member of the pair is (typically) a bird or a mammal that is somehow different from the rest of its class, thus deserves an explicit clarification, such as a whale being a mammal. For some reason, insects also received a fairly large amount of description and thus were easy to contrast with other animals.

The similar pairs are: ape/monkey, bear/panda, bee/moth, beetle/butterfly, beetle/cricket, beetle/fly, bug/caterpillar, butterfly/cricket, butterfly/fly, camel/giraffe, caterpillar/cricket, caterpillar/moth, cricket/fly, donkey/zebra, eagle/hawk, fox/wolf, goose/turkey, horse/pony, leopard/lion, lion/tiger. Again, (an expected) a pattern can be noticed here: those animals that received a lot of similar descriptions are being selected.

There were 7 animals or categories in the dictionary that were used in the is-a relations other than to indicate an offspring of an animal. These categories were: animal, insect, bird, fish, reptile, cat, and horse. Mammal got an entry in the dictionary but was not used in any of the descriptions. We excluded entries that indicated *a young animal*, such *a kitten is a young cat*. We calculated the mean and standard deviation of each animal relative to the above 7 categories using the *hier* metric described above. We assumed that if an entry had a description that X is Y, and hier(X, Y) was lower than the mean for that overall category, the definition should be questioned and should not be used for hierarchy construction. The following entries were so affected: bat is-a animal, crab is-a animal, goat is-a animal, hippo is-a animal, sheep is-a animal, whale is-a animal, ostrich is-a bird, tiger is-a cat, lion is-a cat.

There are several explanations for the results: cat is defined as a domestic animal, and thus, of course, cannot be a parent of wild animals. Crab has more features that puts it next to fish, and so does whale, including the description of the habitat. Hippo is mostly described swimming in lakes and rivers. Bat is similar in its description to a bird. Goat and sheep created a puzzle for us. However, we considered it to be a success to have only 2 problematic entries.

Interestingly, contradicting the dictionary, the metric suggests that donkey and zebra should be types of horse; dog and hamster should be types of cat. These entries suggest that there is not enough differentiation between the affected animals for them to be correctly classified.

We therefore wanted to see if the ratio changes when the omitted common-sense knowledge is added to the descriptions and if some puzzling results are corrected. The common-sense knowledge consisted of a number of additional animal properties, explicitly stated in some descriptions but omitted from others, with clearly implied values, so we added that information directly to the ontology as it emerged from the descriptions. The addition to common-sense knowledge solved the hierarchy problem of animal in the previous experiment not being an animal, and did not introduce any additional problems.

The distribution of the resulting similarity is shown in Figure 2. As seen there, listing results that are 3 standard deviations away from the mean proved to be impractical, although that was done in the first experiment,. Thus, the results below reflect the same number of dissimilar pairs as the first experiment: ape/duck, ape/swan, duck/snail, crocodile/snail, chicken/snail, crow/snail, alligator/snail, ape/goose, beaver/snail, eagle/snail, hawk/snail, goose/monkey, fox/snail, hamster/snail, duck/monkey, ape/snail, deer/swan, ape/crab, goose/snail, camel/snail, jellyfish/monkey, ape/penguin, bear/snail, goat/snail. As with previous results, there is a concept that is most dissimilar to others (snail), and the dissimilarity looks plausible (all below 0).

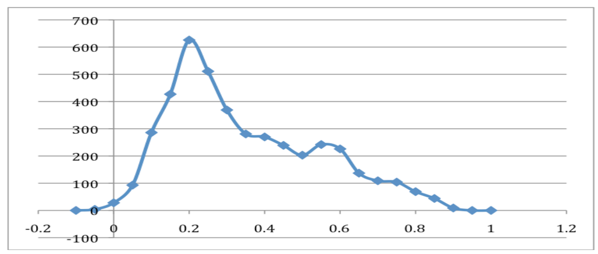

Figure 2. Distribution of similarity of animals

The pairs that are most similar are again several insects and birds, as well as eagle/hawk, crab/lobster, hippo/rhinoceros, bull/cow, dolphin/whale, horse/pony, cow/sheep, cow/pony, cow/sheep, frog/toad, pig/pony, lion/tiger, mouse/rat, jellyfish/octopus, donkey/zebra, spider/worm. As expected, the similarity results look (more?) reasonable with common-sense knowledge (Figure 3).

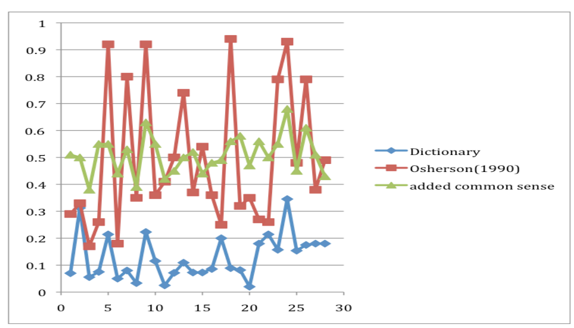

Figure 3: Pair-wise similarity of 7 animals.

While the most similar and least similar results look good, the middle section will need to be improved. Figure 4 shows pair-wise comparisons of perceived similarity between cow, dolphin, elephant, horse, mouse, rhino, seal, and squirrel. We anticipate these results to improve when weights are added to properties.

Conclusion

We have demonstrated, on a very limited corpus of animal descriptions intended for a very young audience, that it is possible to detect semantic structure in natural language descriptions as well as pointing to flawed descriptions. The results improve with the addition of common-sense knowledge omitted from but implied by the descriptions. The specific material, from a children's dictionary that was designed to limit the amount of world knowledge that the young reader could be counted on contributing, helped us delimit the common-sense knowledge. It is clear, of course, that this method of defining this elusive resource is not useful outside of this artificially restricted environment but the convenient handle to it was too tempting to resist.

References

AHFD (2007). The American Heritage First Dictionary. Boston-New York: Houghton Mifflin.

Allen, J. F., Swift, M., & de Beaumont, W. (2008). Deep semantic analysis of text. *Proc. Step '08*. Venice.

Bar-Hillel, Y. (1954). Indexical expressions. *Mind* 63, 359-379. Reprinted in his *Aspects of Language*, Jerusalem: Magnes, 1970, 69-88.

Barwise, J., & Perry J. (1983). *Situations and Attitudes*. Cambridge, MA: MIT Press-Bradford.

Byrne, J., Dale-Tunnicliffe, S., Patrick, T., & Grace, M. (2010). Children's understanding of animals in their everyday life in the UK and USA. *Proceedings of the 7th Conference of European Researchers in Didactics of Biology ERIDOB*.

Carey, S. (1985). *Conceptual Change in Childhood.* Cambridge, MA: MIT Press.

Charniak, E. 1972. Toward a Model of Children's Story Comprehension. Artificial Intelligence Technical Report Number 266, Department of Computer Science, Massachusettes Institute of Technology, Cambridge, MA.

Clark, P., Fellbaum, C., & Hobbs, J. (2008). Using and extending WordNet to support question answering. *Proceedings of the 4th Global WordNet Conference*, 111-119.

Giora, R. (2003). *On Our Mind: Salience, Context, and Figurative Language*. New York: Oxford University Press.

Gordon, J. M., & Schubert L. K.. (2010). "Quantificational Sharpening of Commonsense Knowledge" , in: Havasi et al., 27-32.

Havasi, C., Lenat, D. B., & Van Durme, B. (Eds). (2010). Commonsense Knowledge: Papers from the AAAI Fall Symposium. Menlo Park, CA: AAAI Press, 2010.

Hempelmann, C. F., Taylor, J. M., & Raskin, V. (2010). Application-guided Ontological Engineering. *Proceedings of the International Conference on Artificial Intelligence,* Las Vegas, NE.

Kemp, C., Tanenbaum, J. B., Griffiths, T. L., Yamada, T, and Ueda, N. 2006. Learning systems of concepts with an infinite relational model. *Proc AAAI-06.*

Lenat, D. B. (1990) CYC: Toward programs with common sense, *Communications of the ACM* 33:8, 30–49.

Lewis, D. (1972). General semantics. In Davidson, D. and Harman, G. (Eds.), *Semantics of Natural Language*, Dordrecht - Boston: Reidel.

Liu, H., & Singh, P. (2004). ConceptNet—a practical commonsense reasoning tool-kit, BT Technology Journal, 22:4. 211-226.

McCarthy, J. (1959). Programs with common sense. *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, 75–91. London: Her Majesty's Stationary Office.

Nirenburg, S., & Raskin, V. 2004. *Ontological Semantics*. Cambridge, MA: MIT Press

Opfer, J. E. & Siegler, R. S. (2004). Revisiting preschoolers' *living things* concept: A microgenetic analysis of of conceptual change in basic biology. *Cognitive Psychology* 49, 301-332.

Osherson, D. N., Wilkie, O., Smith, E. E., López, A., and Shafir, E.1990. Category-based induction. Psychological Review, vol. 97, no. 2, 185-200.

Perfors, A., Kemp, C., Tenenbaum, J. B. 2005. Modeling the acquisition of domain structure and feature understanding. *Proc. CogSci-05.*

Prokop, P., Prokop, M., Dale-Tunnicliffe, S., & Diran, C. (2007). Children's ideas of animals' internal structures. *Journal of Biological Education,* 41-2, 62-67.

Raskin, V., Hempelmann, C. F., & Taylor , J. M. (2010). Guessing vs. knowing: The two approaches to semantics in natural language processing, *Annual International Conference Dialogue 2010*, 642-650, Bekasovo (Moscow), Russia.

Taylor, J. M., & Raskin, V. (2011). Understanding the unknown: Unattested input processing in natural language, *Proc. FUZZ-IEEE-11,* Taipei, Taiwan.

Taylor, J. M., Hempelmann, C. F., & Raskin, V. (2010a). On an automatic acquisition toolbox for ontologies and lexicons in Ontological Semantics, *International Conference on Artificial Intelligence,* Las Vegas, NE.

Taylor, J. M., Raskin, V. & Hempelmann, C. F. (2011a). From disambiguation failures to common-sense knowledge acquisition: A day in the life of an ontological semantic system. In Gilof, R., (Ed.), *Proceedings of the Web Intelligence and Intelligent Agent Technology Conference*. IEEE.

Taylor, J. M., Raskin, V., & Hempelmann, C. F. (2011b). Towards computational guessing of unknown word meanings: The Ontological Semantic approach, *Proc of CogSci-11*, Boston, MA.

Tunnicliffe S. D., Boulter, C., & Reiss, M. (2007). Pigeon – friend or foe? Children's understanding of an everyday animal. *Proceedings of the British Educational Research Association Annual Conference*.

Van Durme, B., Michalak, P., & Schubert, L.K. (2009). Deriving Generalized Knowledge from Corpora Using WordNet Abstraction. *Proceedings of EACL*, 808-816.