OMCSNet: A Commonsense Inference Toolkit

Hugo Liu and Push Singh
MIT Media Laboratory
20 Ames St., Bldg. E15
Cambridge, MA 02139 USA
{hugo, push}@media.mit.edu

Abstract

Imparting commonsense knowledge to computers enables a new class of intelligent applications better equipped to make sense of the everyday world and assist people with everyday tasks. While previous attempts have been made to acquire and structure commonsense knowledge, they have either been inadequate in capturing the breadth of knowledge needed for the enterprise, or their complicated representation schemes have made them difficult to incorporate into applications.

In this paper we describe OMCSNet, a freely available commonsense knowledge base that at once possesses great breadth of knowledge and that can be easily incorporated into applications. Built from the Open Mind Common Sense corpus, which acquires commonsense knowledge from a web-based community of instructors, OMCSNet is a semantic network of 280,000 items of commonsense knowledge, and a set of tools for making inferences using this knowledge. We describe the structure and contents of OMCSNet and its associated inference toolkit, review applications that have incorporated it, and evaluate and analyze this resource.

1 Introduction

The evolution of intelligent software is quickly reaching the point where purely statistical methods or narrow sets of domain-specific rules no longer suffice. There is an increasing demand for the breadth of knowledge about people and the everyday world covered only by a commonsense knowledgebase. While the pursuit of imparting commonsense knowledge to computers is as old as AI itself, progress has been slow in the area of building adequate databases of commonsense knowledge about everyday life.

The problem is largely one of scale, for it has been estimated that the scope of commonsense may involve many tens of millions of pieces of knowledge. The magnitude of this task has discouraged most artificial intelligence researchers from attacking the problem directly. However, it is increasingly clear that the lack of publicly available commonsense resources in the AI community is stifling innovation toward more intelligent computer systems.

In order to encourage innovation in research and enable applications that require large-scale commonsense knowledge bases, we built **OMCSNet**, a semantic network of 280,000 items of commonsense knowledge. An excerpt of OMCSNet is shown in Figure 1. Our aim was to create a large-scale machine-readable resource structured as an easy-to-use semantic network representation like WordNet [Fellbaum, 1998] and MindNet [Richardson *et al.*, 1998], yet whose contents reflect the broader range of world knowledge characteristic of commonsense as in Cyc [Lenat, 1995]. While far from an ideal commonsense inference system, OMCSNet has nonetheless offered the knowledge and inference mechanisms to support plausible commonsense inference in a variety of ap-

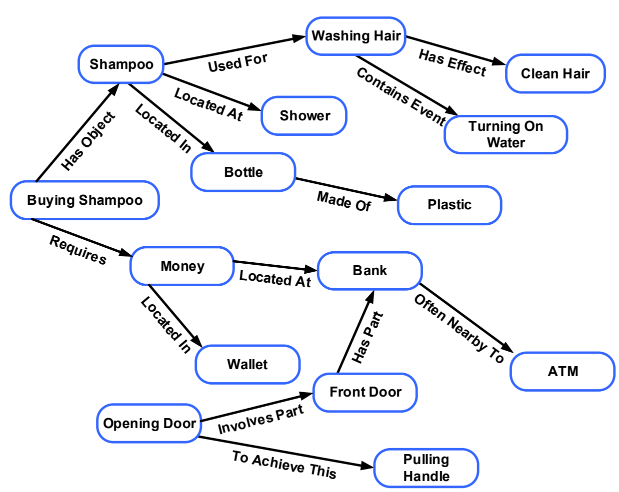

Figure 1. An excerpt from OMCSNet

plications.

This paper is structured as follows. First, we discuss how OMCSNet was built, how it is structured, and the nature of its contents. Second, we present the OMCSNet inference toolkit and briefly review applications that have incorporated it. Third, we evaluate several aspects of the knowledge and the inference toolkit, and compare it to several other large-scale semantic knowledge bases. We conclude with a discussion of the potential impact of this resource on the AI community at large, and explore directions for future work.

2 OMCSNet

In this section, we first explain the origins of OMCSNet in the Open Mind Commonsense corpus; then we demonstrate how knowledge is extracted to produce the semantic network; and third, we describe the structure and semantic content of the network.

2.1 Building OMCSNet

We built OMCSNet in a unique way. Three years ago we built the Open Mind Commonsense (OMCS) web site [Singh *et al.* 2002], a collection of 30 different activities each of which elicits a different type of commonsense knowledge—simple assertions, descriptions of typical situations, stories describing ordinary activities and actions, and so forth. Since then we have gathered nearly 500,000 items of commonsense knowledge from over 10,000 contributors from around the world, many with no special training in computer science. The OMCS corpus now consists of a tremendous range of different types of commonsense knowledge, expressed in natural language.

The earliest applications of the OMCS corpus made use of its knowledge not directly but by first extracting into semantic networks only the types of knowledge they needed. For example, the ARIA photo retrieval system [Lieberman & Liu, 2002a] extracted taxonomic, spatial, functional, causal, and emotional knowledge to improve information retrieval. This suggested to us a new approach to building a commonsense knowledgebase. Rather than directly engineering the knowledge structures used by the reasoning system, as is done in Cyc, we instead encourage people to provide information clearly in natural language and then extract from that more usable knowledge representations. We were inspired by the fact that there had been significant progress in the area of information extraction from text in recent years, due to improvements in broad-coverage parsing [Cardie, 1997]. A number of systems are able to successfully extract facts, conceptual relations, and even complex events from text.

OMCSNet is produced by an automatic process, which applies a set of ‘commonsense extraction rules’ to the OMCS corpus. A pattern matching parser uses 40 mapping rules to easily parse semi-structured sentences into predicate relations and arguments which are short fragments of English. These arguments are then normalized using natural language techniques (stripped of stop words, lemmatized), and are *massaged* into one of many standard syntactic forms. To account for richer concepts which are more than words, we created three categories of concepts: Noun Phrases (things, places, people), Attributes (modifiers), and Activity Phrases (actions and actions compounded with a noun phrase or prepositional phrase, e.g.: “turn on water,” “wash hair.”). A small part-of-speech tag –driven grammar filters out non-compliant text fragments and massages the rest to take one of these standard syntactic forms. When all is done, the cleaned relations and arguments are linked together into the OMCSNet semantic network.

2.3 Contents of OMCSNet

At present OMCSNet consists of the 20 types of binary relations shown below in Table 1. These relations were chosen because the original OMCS corpus was built largely through its users filling in the blanks of templates like ‘a hammer is for _____’. Thus the relations we chose to extract largely reflect the original choice of templates used on the OMCS web site.

Relation Type	Semantic Relation
Things	IsA, HasProperty, PartOf, MadeOf
Events	SubEventOf, NextEvent, FirstStepOf, LastStepOf
Actions	Requires, HasEffect, ResultsInWant, HasAbility
Spatial	OftenNear, LocationOf, CityInLocality
Goals	DoesWant, DoesNotWant, MotivatedBy
Functions	UsedInLocation, HasFunction
Generic	ConceptuallyRelatedTo

Table 1. Semantic Relation Types currently in OMCSNet

The OMCSNet Browser Tool can be used to browse the contents of OMCSNet by searching for concepts and following semantic links. A picture of this tool is shown in Figure 2.

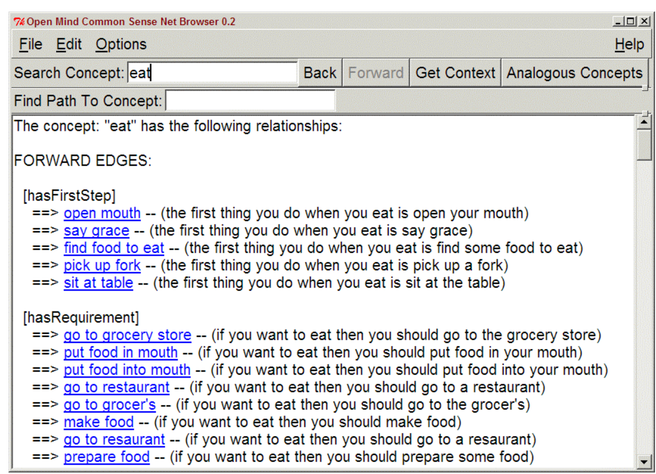

Figure 2. The OMCSNet Browser Tool

3 OMCSNet Inference Toolkit

To assist in using OMCSNet in various types of inference, we built a small but growing set of tools to help researchers and application developers maintain a high-level, task-driven view of commonsense. In the following subsections, we describe some of the more basic tools.

'Fuzzy' Inference. So far we have presented OMCSNet as a fairly straightforward semantic network, and so one might ask the question why an inference toolkit might even be necessary when conventional semantic network graph traversal techniques should suffice. The answer lies in the structure of the nodes, and in the peculiarity of commonsense knowledge.

In the previous section we presented several types of nodes including Noun Phrases, Attributes, and Activity Phrases. These nodes can either be first-order, i.e. simple words and phrases, or second-order, such as "turn on water." Second order nodes are essentially fragments of English following a particular part-of-speech pattern. Maintaining the representation in English saves us from having to map into and out of a special ontology, which would greatly increase the complexity and difficulty-of-use of the system; it also maintains the nuances of the concept. Practically, however, we may want the concepts "buy food" and "purchase food" to be treated as the same concept.

To accomplish this, the inference mechanism accompanying OMCSNet can perform such *fuzzy* conceptual bindings using a simple semantic distance heuristic (e.g. "buy food" and "purchase food" are commensurate if a synonym relation holds between "buy" and "purchase.") Another useful approximate matching heuristic is to compare normalized morphologies produced by lemmatizing words. Using these approximate concept bindings, we can perform 'fuzzy' inference over the network.

Context Determination. One task useful across many natural language applications is determining the context around a concept or around the intersection of several concepts. The context determination tool enables this by performing spreading activation to discover concepts in the *semantic neighborhood*. For example, OMCSNet produced the following top concepts in the neighborhood of the concept "living room," (Table 2). Percentages indicate confidence of semantic connectedness.

house (100.0%)	kitchen (94.7%)	library (88.1%)
home (99.6%)	table (93.2%)	roof (87.4%)
apartment (98.4%)	bed (92.4%)	couch (87.4%)
bedroom (97.1%)	town (92.0%)	chair (86.9%)
building (96.5%)	office (91.0%)	rug (86.5%)
floor (95.3%)	next door (89.3%)	comfortable (84.1%)
room (95.2%)	wall (88.8%)	coffee table (83.2%)

Table 2. Concepts in the semantic neighborhood of "living room" (over all relations without conceptual bias)

Concepts connected to "living room" through any relation were included in the context. However, we may, for example, only be interested in specific relations. If we had specified the relation "HasFunction", the context search would return results like "entertain guests," "comfortable," and "watch television." In other cases we may desire to bias the context of "living room" with another concept, e.g., "store." The output is the context of "living room" with respect to the concept "store" and returns results like "furniture," "furniture store," and "Ikea."

Analogical Inference. Knowledge about particular concepts is occasionally spotty. For example, the system may know "Requires(car, gas)" but not "Requires(motorcycle, gas)". Such relationships may be produced using analogical inference. For example, by employing structure-mapping methods [Falkenhainer *et al*., 1989]. In the present toolkit, we are already able to make some simple conceptual analogies using structure-mapping, producing results like the following (Figure 3):

```
car is like motorcycle because both:
  ==[IsA]==> vehicle type
  ==[HasFunction]==> transportation
  ==[HasProperty]==> fast
```

Figure 3. An example of conceptual analogy over OMCSNet using structure-mapping.

4 How OMCSNet is being used

Early versions of the OMCSNet tools are being put to use in a variety of prototype applications, each of which uses commonsense differently. None of them actually does 'general purpose' commonsense reasoning. While each makes use of a broad range of commonsense knowledge, each makes use of it in a particular way by performing only certain types of inferences.

Retrieving event-subevent structure. It is sometimes useful to collect together all the knowledge that is relevant to some particular class of activity or event. For example the Cinematic Common Sense project makes use of commonsense knowledge about event-subevent structure in OMCSNet to make suitable shot suggestions at common events like birthdays and marathons [Barry & Davenport, 2002]. For the topic 'getting ready for a marathon', the subevents gathered might include: putting on your running shoes, picking up your number, and getting in your place at the starting line.

Goal recognition and planning. The search engines described in Singh [2002] and Liu *et al.* [2002] exploit commonsense knowledge about typical human goals to infer the real goal of the user from their search query. For example, the search 'my cat is sick' leads to the system inferring that 'I want my cat to be healthy' because people care about their pets and they want things they care about to be healthy. Furthermore, these search engines can make use of knowledge about actions and their effects to engage in a simple form of planning. After inferring the user's true intention, they look for a way to achieve it. In this case, if you want something to be healthy you can take it to a doctor, or in the case of an animal, a veterinarian.

Temporal projection. The MakeBelieve storytelling system [Liu & Singh, 2002] makes use of the knowledge of temporal and causal relationships between events in order to guess what is likely to happen next. Using this knowledge it generates stories such as: *David fell off his bike. David scraped his knee. David cried like a baby. David was laughed at. David decided to get revenge. David hurt people.*

Particular consequences of broad classes of actions. Empathy Buddy senses the affect in passages of text [Liu et al., 2003]. It predicts those consequences of actions and events that have some emotional significance. This can be done by chaining backwards from knowledge about desirable and undesirable states. For example, if being out of work is undesirable, and being fired causes to be to be out of work, then the passing 'I was fired from work today' can be sensed as undesirable.

Specific facts about particular things. Some of OMCSNet is specific facts like "the Golden Gate Bridge is located in San Francisco", or that "a PowerBook is a kind of laptop computer." The ARIA e-mail client and photo retrieval system [Liu & Lieberman, 2002] can reason that an e-mail that mentions that "I saw the Golden Gate Bridge" meant that I was in San Francisco at the time, and proactively retrieves photos taken in San Francisco for the user to insert into the e-mail.

Conceptual association. OMCSNet can be used to supply associated concepts. The Globuddy program [Various Authors, 2003] uses OMCSNet to retrieve knowledge about events, actions, objects, and other concepts related to a given situation, to make a custom phrasebook of concepts you might wish to have translations for in that situation. For example, if you are arrested, it will give you a few pages translating words like 'lawyer', 'going to prison', 'find a lawyer', and so forth.

5 Evaluation

The original OMCS corpus was previously evaluated by Singh *et al.* [2002]. Human judges evaluated a sample of the corpus and rated 75% of items as largely true, 82% as largely objective, 85% as largely making sense, and 84% as knowledge someone would have by high school.

We performed two further analyses of OMCSNet: a qualitative study (human judges) and a quantitative analysis. However, perhaps the most compelling evaluations are indirect. OMCS and OMCSNet have been used to measurably improve the behavior of intelligent agents. In the previous section we briefly reviewed some OMCS- and OMCSNet- enabled agents. For brevity, we refer the reader to each application's respective evaluations (see each application's corresponding paper).

A Qualitative Study of OMCSNet. We conducted an experiment with five human judges and asked each judge to rate 100 concepts in OMCSNet. 10 concepts were common to all judges (for correlational analysis), 90 were of their choice. If a concept produced no results, they were asked to duly note that and try another concept. Concepts were judged along these 2 dimensions, each on a Likert 1 (strongly disagree) to 5 (strongly agree) scale:

1) Results for this concept are fairly comprehensive.
2) Results for this concept include incorrect knowledge, nonsensical data, or non-commonsense information.

To account for inter-judge agreement, we normalized scores using the 10 common concepts, and produced the re-centered aggregate results shown below in Table 3.

	Mean Score	Std. Dev.	Std. Err.
Comprehensiveness	3.40 / 5.00	1.24	1.58
Noisiness	1.24 / 5.00	0.99	1.05
% Concepts attempted, that were not in KB	11.3%	6.07%	0.37%

Table 3. Measure of quality of OMCSNet.

These results can be interpreted as follows. Judgment of comprehensiveness of knowledge in OMCSNet on average, was *several relevant concepts,* but varied significantly from *a few concepts* to *almost all of the concepts.* Noisiness was *little noise* on average, and did not vary much. % of KB misses was very consistently 11%. We consider these to be very optimistic results. Comprehensiveness was moderate but varied a lot indicating still spotty coverage, which we hope this will improve as OMCS grows. Noisiness was surprisingly low, lending support to the idea that a relatively clean KB can be elicited from public acquisition. % of KB misses was more than tolerable considering that OMCSNet has only 80,000 concepts—a fraction of that possessed by people.

A Quantitative Analysis of OMCSNet. 100 salient concepts already in OMCSNet were selected by the judges for each contextual "domain" as typifying that domain (for example, the domain of "everyday", concepts includes "wake up", "eat breakfast", "shower", "go to work", "prepare meal", "eat food", etc.). Concepts included appropriate distributions of concept types, i.e. people, places, things, actions, and activities. Branching factor indicates the number of relations for each node (density of knowledge). Standard deviation illustrates unevenness of knowledge. The intra-set branching factor and standard deviations indicate density and unevenness *within* each domain. Results are shown in Table 4.

	Overall KB	Jobs	Family	Every-day	Trips
Branching Factor	3.48	59.7	98.5	40.1	34.9
Standard Dev.	21.5	78.6	169	38.5	38.3
Intra-set B.F.		4.06	8.83	2.2	1.7
Intra-set Std. Dev.		4.17	9.75	2.35	2.05

Table 4. Coverage density and distribution in 4 domains.

These results show that although there is a lot of knowledge about these common domains, there is also an enormous variation of coverage. A review of the histogram of results (not shown) indicates a camel distribution—a concept possessed either a lot of knowledge (>100) or not much (<5). We postulate that structure of the semantic network consists of mainly dense "hub" nodes (possibly due to word-sense collision) and some outlying spoke nodes. From the intra-set results, we postulate that knowledge is not as clustered around domains as we had expected. This is an interesting result because it suggests that artificial clustering of domains prevalently practiced in AI may not work for commonsense!

6 Large-Scale Semantic KBs

In this section we compare OMCSNet with several other existing large-scale semantic knowledge bases.

Cyc. The Cyc project (Lenat 1995) is the most prominent large-scale effort to build a commonsense knowledge base. A major difference between Cyc and OMCSNet is in the choice of knowledge representation. Knowledge in Cyc is represented in a rich logical language called CycL. OMCSNet, on the other hand, explores an alternative representation grounded in structured English fragments and a limited set of predicate relations. So OMCSNet loosely resembles predicate logic over fragments of English. OMCSNet's semantic network is a much simpler and less expressive knowledge representation scheme than CycL, and as a result OMCSNet cannot represent many important types of commonsense knowledge. While not as formal as CycL, we nonetheless believe that a broad range of applications still stand to benefit from such a knowledge base.

From a practical perspective, another important difference is that the Cyc knowledge base is at present proprietary and inaccessible as a community resource, whereas both the OMCS corpus and OMCSNet are freely available resources. However, recently the developers of Cyc have released OpenCyc, a publicly available version of Cyc that includes its inference engine and Cyc's upper level ontology.

ThoughtTreasure. With on the order of 100,000 items of commonsense knowledge, ThoughtTreasure (TT) was built by researcher Erik Mueller to investigate the story understanding task (Mueller, 1998). TT represents commonsense knowledge in a variety of ways, including simple assertions, frames, scripts, and spatial occupancy arrays. The knowledge in TT is well-mapped onto natural language, for every concept has an associated lexical item, and the TT system itself includes a substantial natural language parsing and generation component. By comparison, knowledge in OMCSNet is completely assertional (although the OMCS corpus itself contains other types of knowledge that were not included in OMCSNet), and its representation is rooted in semi-structured English fragments.

WordNet. Arguably the most widely used machine-readable semantic resource in the artificial intelligence and computational linguistic communities, WordNet was not intended as a commonsense resource per se, but rather as a large lexical database of English concepts (simple words and collocations). The scope of WordNet encompasses on the order of 100,000 concepts, connected by 100,000 nymic relations of hypernymy (is-a), hyponymy (a-kind-of), synonymy, antonymy, and meronymy (part-of). It is attractive as a commonsense resource because its hierarchical system of concepts captures some basic (but limited) relationships between concepts in the everyday world, and is comprehensive enough to have wide application.

WordNet's popularity with researchers and developers illustrates the two communities' thirst for semantic knowledge bases. Its representational simplicity (all binary relations) and its being rooted in plain English (no complex representational language to map into or out of) lends it an ease of use and integration into applications that has also promoted adoption. We feel that OMCSNet, with a comparable knowledge representation but offering more diverse semantic content, will also help to address the knowledge needs of the communities and foster innovation that would not be possible otherwise.

OMCSNet differs from WordNet in a few important ways. First, concepts in OMCSNet are not sense disambiguated as in WordNet, though it is possible to introduce a statistical notion of "sense" by clustering conceptual nodes in a graph. Second, concepts in WordNet are organized into syntactic categories of nouns, verbs, adjectives, and adverbs and are usually one word or a collocation with one head word; in contrast, concepts in OMCSNet contain a variety of semantic categories like things, people, properties, actions, activities, and events, and may contain many hyperlexical concepts (e.g. "buy groceries") in addition to lexical ones. Third, relations in WordNet are primarily hierarchical and are limited in the relationships they can express; OMCSNet presently uses 20 relations including temporal, spatial, causal, and functional relations, which are arguably more useful for commonsense reasoning problems.

MindNet. OMCSNet and MindNet follow a very similar approach. Also based on the premise that large, useful semantic networks can be extracted from natural language text corpora, the MindNet project mines reference materials like dictionaries using broad-coverage parsing techniques to populate a semantic network with named relations. The two semantic networks have comparable numbers of named semantic relations, and go beyond basic WordNet nymic relations, which are largely hierarchical. However, there are several pointed differences.

First, MindNet is fundamentally a *lexical* knowledge base—concepts that are words. This reflects the fact that they are parsing primarily lexical resources including Longman's Dictionary of Common English (LDOCE) and American Heritage Dictionary, 3rd edition (AHD3); in addition, imperfect broad coverage parsing over unstructured text (dictionary definitions are unstructured) makes it hard to parse relationships between entities much larger than individual words. Because of its knowledge source, OMCSNet's concept nodes are often hyperlexical (second order nodes), including English fragments such as activity phrases (e.g. "wash hair", "brush teeth"), and concept phrases (e.g. "automatic teller machine"). For the same reason, MindNet's relations primarily describe lexical-level properties such as Part, Possessor, Material, Source, etc. As a result, non-lexical commonsense not resembling dictionary definitions is harder to express in the MindNet formalism, e.g. "eating a lot of food will make you less hungry."

Second, MindNet relies on dictionary corpora, and dictionary definitions and wording are often not very representative of the *practical* and *everyday* meaning of concepts. Mined from dictionaries, MindNet will provide only one or two definitions of each concept, while we maintain a plurality of ways of representing a concept's meaning, and a plurality of different ways to phrase a definition. Mining of dictionaries and reference resources may be useful for acquiring a small subset of denotational, lexical commonsense, but ultimately a large part of commonsense is not written in existing references.

Relative sizes of Knowledgebases. Table 4 compares the sizes of these five large-scale semantic knowledgebases. The size of Cyc is on the order of 1.5 million assertions, though we caution that numbers given throughout this section are specific to each project's knowledge representation and therefore they should be compared with caution.

Name	Concepts	ako/isa	part-of	Other
Cyc	30,000	25%	35%	40%
ThoughtTreasure	27,093	28,818	666	21,821
WordNet 1.6	99,642	78,446	19,441	42,700
MindNet	45,000	47,000	14,100	32,900
OMCSNet	81,430	45,382	5,208	151,692

Table 4. The relative size of knowledgebases. Adapted with permission from Mueller (2002)).

7 Extending OMCSNet

We are presently extending OMCSNet in several directions. First, we would like to disambiguate the senses of the concepts in OMCSnet. The Open Mind Word Expert web site (Chklovski and Mihalcea, 2002) allows users to disambiguate the senses of the words in the OMCS corpus, and we are looking into making use of the data they are collecting to build a disambiguated OMCSNet. Second, the current set of 20 relation types in OMCSNet is small compared to the wide array of assertion types that exist in the OMCS corpus. We wish to employ a broad coverage parser that can extract a wider range of knowledge from the corpus. Third, we are developing a special version of the OMCS web site that focuses specifically on further growing the OMCSNet knowledge base, including special activities for elaborating, validating, repairing items of knowledge.

Conclusions

OMCSNet is presently the largest freely available database of commonsense knowledge. It comes with a browser and a preliminary set of inference tools, and is being used in a number of applications. While the contents of the knowledgebase are still spotty in comparison to what people know, our analysis has shown it to be surprisingly clean, and it has proven more than large enough to enable experimenting with entirely new kinds of interactive applications with common sense.

References

[Barry & Davenport, 2002] Barry, B. & Davenport. G. (2002). *Why Common Sense for Video Production?* (Interactive Cinema Technical Report #02-01). Media Lab, MIT.

[Cardie, 1997] Cardie, C. (1997). Empirical Methods in Information Extraction, *AI Magazine*, 65-79.

[Chklovski & Mihalcea, 2002] Chklovski, T. and Mihalcea, R. (2002). Building a Sense Tagged Corpus with Open Mind Word Expert. In *Proceedings of the Workshop on "Word Sense Disambiguation: Recent Successes and Future Directions", ACL 2002.*

[Falkenhainer *et al.*, 1990] Falkenhainer, B., Forbus, K.D. and Gentner, D. (1990). The structure-mapping engine: algorithm and examples. *Artificial Intelligence*, 41:1-63.

[Fellbaum, 1998] Fellbaum, Christiane. (Ed.). (1998). *WordNet: An electronic lexical database*. Cambridge, MA: MIT Press.

[Lenat, 1995] Lenat, D. B. (1995). CYC: A large-scale investment in knowledge infrastructure. CACM 38(11): 33-38.

[Lieberman *et al.,* 2001] Lieberman, H., Rosenzweig E., Singh, P., (2001). Aria: An Agent For Annotating And Retrieving Images, *IEEE Computer*, July 2001, 57-61.

[Liu & Lieberman, 2002] Liu, H., Lieberman, H. (2002). Robust photo retrieval using world semantics. *Proceedings of the LREC2002 Workshop: Creating and Using Semantics*. Las Palmas, Canary Islands.

[Liu, Lieberman & Selker, 2002] Liu, H., Lieberman, H., Selker, T. (2002). GOOSE: A Goal-Oriented Search Engine With Commonsense. *Proceedings of AH2002.* Malaga, Spain.

[Liu *et al.,* 2003] Liu, H., Lieberman, H., Selker, T. (2003). A Model of Textual Affect Sensing using Real-World Knowledge. In *Proceedings of IUI 2003.* Miami, Florida.

[Liu & Singh, 2002] Liu, H., Singh, P. (2002). MAKEBELIEVE: Using Commonsense to Generate Stories. In *Proceedings of AAAI-02*. Edmonton, Canada.

[Mueller, 2000] Mueller, E. (2000). ThoughtTreasure: A natural language/commonsense platform. Retrieved from http://www.signiform.com/tt/htm/overview.htm on 1/9/03.

[Richardson *et al.*, 1998] Richardson, S. D., Dolan, B., and Vanderwende, L. (1998). MindNet: Acquiring and structuring semantinc information from text. In *COLING-ACL'98*.

[Singh *et al.,* 2002] Singh, Push, Lin, Thomas, Mueller, Erik T., Lim, Grace, Perkins, Travell, & Zhu, Wan Li (2002). Open Mind Common Sense: Knowledge acquisition from the general public. In *Proceedings of ODBASE'02*. Lecture Notes in Computer Science. Heidelberg: Springer-Verlag.

[Various Authors, 2003] Various Authors (2003). Common Sense Reasoning for Interactive Applications Projects Page. Retrieved from http://www.media.mit.edu/~lieber/Teaching/Common-Sense-Course/Projects/Projects-Intro.html on 1/9/03.