

Machine Reading as a Process of Partial Question-Answering

Peter Clark and Phil Harrison

Boeing Research & Technology
The Boeing Company, PO Box 3707, Seattle, WA 98124, USA
{peter.e.clark, philip.harrison}@boeing.com

Abstract

This paper explores the close relationship between question answering and machine reading, and how the active use of reasoning to answer (and in the process, disambiguate) questions can also be applied to reading declarative texts, where a substantial proportion of the text's contents is already known to (represented in) the system. In question answering, a question may be ambiguous, and it may only be in the process of trying to answer it that the "right" way to disambiguate it becomes apparent. Similarly in machine reading, a text may be ambiguous, and may require some process to relate it to what is already known. Our conjecture in this paper is that these two processes are similar, and that we can modify a question answering tool to help "read" new text that augments existing system knowledge. Specifically, interpreting a new text T can be recast as trying to answer, or partially answer, the question "Is it true that T ?", resulting in both appropriate disambiguation and connection of T to existing knowledge. Some preliminary investigation suggests this might be useful for proposing knowledge base extensions, extracted from text, to a knowledge engineer.

1 Introduction

Machine reading is not just a task of language processing, but an active interplay between knowledge and language; Prior knowledge should guide interpretation of new text, and new interpretations should augment that prior knowledge. Such interaction is essential if ambiguities in language are to be resolved "correctly" (with respect to what is

known), and if the resulting interpretations are to be integrated with existing knowledge. The main insight of this paper is that this interaction is similar to that required for knowledge-based question answering, which also requires searching a knowledge base (KB) for a valid interpretation of the question. In our earlier work on question answering (Clark and Harrison, 2010), we found that some disambiguation decisions for question interpretation could be deferred, to be resolved during question answering, guided by what was found in the KB. In this paper, we show how a similar approach can be applied to interpreting declarative text, so that a similar interplay between language and knowledge is achieved.

"Machine reading" itself is a loosely-defined notion, ranging from extracting selective facts to constructing complex, inference-supporting representations of text. One approach for selective extraction is the use of semantic templates ("scripts", "frames") to provide a set of roles (slots) and constraints on objects playing those roles (fillers) to be expected in text, and might be filled by methods ranging from simply skimming text, e.g., FRUMP (DeJong, 1979), to full language processing, e.g., (Dahlgren et al., 1991). Other work has looked at techniques for learning phrasal patterns likely to contain slot fillers (Riloff, 1996; Sekine, 2006) or contain information semantically similar to a set of seed examples (Carlson et al, 2009).

At the other end of the spectrum, some systems attempt a full understanding of text, i.e., have the ambitious goal of building a complete representation of the text's contents (e.g., Zdrozny 1991, Hobbs et al, 1993). A common thread of these approaches is to search a space of alternative disambiguations and elaborations and select the most

"coherent", based on criteria such as maximizing coreference, minimizing redundancy, and avoiding contradictions. For example, Mulkar et al (2007) search for a set of abductive inferences on the (logical form of the) text that minimizes cost (maximizes coherence) of the result, where an abductive inference might be a word sense or coreference decision with an associated cost. Similarly, Zadrozny and Jensen (1991) search a space of disambiguations when interpreting paragraphs by elaborating each alternative (using dictionary definitions) and selecting the most coherent based on similar criteria. Work on model building is inspiring but also challenging due to the lack of constraint on the final models (even with substantial domain knowledge) and the difficulty of quantifying "coherence".

Our work falls somewhere between these two. We do not use templates for new knowledge, but rather use inference at run-time to identify what is known and thus what to expect that the text might be saying. However, unlike full model building approaches, we assume that the majority of what is being read is already known (represented) in the KB, and thus the reading task is primarily one of recognizing that knowledge in the text, and extending it with any new facts that are encountered. We might term this a "model extension" approach; it corresponds to Feigenbaum's (2003) challenge of, given the representation of a book, have a machine read a second book (about the same topic) and integrate the new knowledge contained in that text.

2 The Problem

Our work is in the context of cell biology, where we have a moderately sized¹, hand-built knowledge base available containing formal representations of biological structures and processes expressed in first-order logic. Our goal is to take paragraphs of text about a topic partially covered by the KB, and identify facts which are already known, facts which are new, and the connections

¹ Specifically, it has detailed representations of entities and processes related to cell division, DNA replication, and protein synthesis, containing approximately 250 domain-specific concepts (built on top of a pre-existing library of approximately 500 domain-general concepts), 120 relations (binary predicates), and approximately 2000 axioms, built as part of Project Halo (Gunning et al., 2010).

Topic: prophase

Input Paragraph:

In the cytoplasm, the mitotic spindle, consisting of microtubules and other proteins, forms between the two pairs of centrioles as they migrate to opposite poles of the cell.

Output Axioms: (expressed in English)

In all prophase events:

- a. The mitotic spindle has parts the microtubule **and the protein.**
- b. The mitotic spindle is created **between the centrioles in the cytoplasm.**
- c. The centrioles move **to the poles.**

Figure 1: The system's behavior, showing known (normal font) and new (bold) facts identified from the text. Note that the output is not just a simple recitation of the input, but a mixture of known and new axioms for the KB.

between the two. An example of the system's output is shown in Figure 1. In the Output Axioms in that Figure, the normal font shows facts that the system has recognized as already known in the KB, while bold font shows new knowledge. It is important to note that the output facts are not just a simple recitation of the input, but have been interpreted in the context of the KB. For example in the input paragraph:

"the mitotic spindle, consisting of microtubules"

has not been interpreted as describing some "consisting" event, but recognized (via use of paraphrases, described later) as referring to the has-part(mitotic-spindle01,microtubules01) element in the representation of prophase in the KB, i.e., denoting a "has part" relationship ((a) in Figure 1). Similarly,

"the spindle forms"

has not been interpreted as an organizing event ("form a team") nor as the spindle doing the forming ("the spindle forms something"), but instead been recognized as the result(create01,mitotic-spindle01) element in the representation of prophase in the KB, i.e., "forms" has been interpreted as this particular creation event in the representation of prophase (b in Figure 1). Doing this requires not just careful language processing; it requires querying the knowledge base to see/infer

what is already known, and using this to guide the interpretation.

This process is similar to that needed for question-answering. Consider giving a question form of (part of) the earlier paragraph to a question-answering system:

- (1) Is it true that the mitotic spindle consists of microtubules?

Again, the phrase "consists of" is ambiguous, and may mean different things in different contexts. However, in the context of question-answering, there is a natural approach to disambiguating this: as the user is asking about what is in the KB, then a natural approach is to query the KB for relationships that hold between the mitotic spindle and microtubules, and see if any are a plausible interpretation of "consists of". If there is one, then it is likely to be the interpretation that the user intended (assuming the user is not being deliberately obscure; we call this a "benevolent user" assumption). If this happens, then the question-answering system can answer "yes"; but more importantly from a machine reading point of view, the system has also correctly disambiguated the original question and located the facts it mentions in the knowledge base as side-effects. It is this process that we want to apply to interpreting declarative texts, with the change that unproven parts should be treated as new assertions, rather than failed queries.

3 Approach

Based on these observations, our approach is to interpret a new text T by treating it as a question to the KB asking whether the facts in T are already known. By attempting to answer this question, the system resolves ambiguity for the known facts (namely, the resolution that leads to them being recognized is preferred). For new facts, the system falls back on more traditional NLP modules, filtered by coarser-grained type constraints. In addition, the identification of known facts in the KB and the connection between the old facts and the new facts provides anchor points for the new facts to be connected to.

To implement this approach, we have used three important features of our question-answering system, here reapplied to the task of text interpretation:

- a. The use of a large database of paraphrases to explore alternative phrasings (hence alternative interpretations) of text;
- b. Deferring word sense and semantic role commitment during initial language processing, to be resolved later based on what is found in the KB;
- c. The use of standard disambiguation techniques to process new facts not located in the KB.

We now summarize these three features, then present the complete algorithm.

3.1 Paraphrases

A relatively recent advance in NLP has been the automatic construction of paraphrase databases (containing phrasal patterns with approximately equivalent meaning), built by finding phrases that occur in distributionally similar contexts (e.g., Dras et al, 2005). To date, paraphrase databases have primarily been exploited for recognizing textual entailment (e.g., Bentivogli et al., 2009). In our work, we take them in a new direction and exploit them for language interpretation.

We use the DIRT paraphrase database (Lin and Pantel, 2001a,b), containing approximately 12 million automatically learned rules of the form:

IF X *relation* Y THEN X *relation'* Y

where *relation* is a path in the dependency tree between constituents X and Y , or equivalently (as we use later) a *chain of clauses*:

$$\{p_0(x_0, x_1), w_1(x_1), \dots, p_{n-1}(x_{n-1}, x_n)\}$$

where p_i is the syntactic relation between (non-prepositional) constituents x_i and x_{i+1} , and w_i is the word used for x_i . An example from DIRT is:

IF X is found in Y THEN X is inside Y

The condition "X is found in Y" can be expressed as the clause chain:

$$\{ \text{object-of}(x, f), \text{"find"}(f), \text{"in"}(f, y) \}$$

We use DIRT to explore alternative interpretations of the text, singling out those that help identify the facts in the text that are already known in the KB.

3.2 Deferred Sense Commitment

Two common challenges for NLP are word sense disambiguation (WSD) and semantic role labeling

(SRL). While there are a number of existing tools for performing these tasks based on the linguistic context (e.g., Toutanova et al., 2008, Erk and Pado, 2006), their performance is only moderate (e.g., Agirre et al, 2007). The problem is accentuated when trying to disambiguate in a way consistent with a particular KB, because there is often a degree of subjectivity in how the knowledge engineer chose to represent the world in that KB (e.g., whether some object is the "agent" or "instrument" or "site" of an activity is to a degree a matter of viewpoint). Trying to create a WSD or SRL module that reliably mimics the knowledge engineer's decision procedure is difficult.

To address this, we defer WSD and SRL commitment during the initial text processing. Instead, these ambiguities are resolved during the subsequent stage of querying the KB to see if (some interpretation of) the text is already known. One can view this as a trivial form of preserving under-specification (eg. Pinkal, 1999) in the initial language processing, where the words themselves denote their possible meanings.

3.3 Interpretation of New Knowledge

Given some text, our system attempts to disambiguate it by searching for (some interpretation of) its statements in the KB. However, this will only disambiguate statements of facts that are already known. For new facts, we fall back on traditional disambiguation methods, using a set of approximately 100 hand-built rules for semantic role labeling, and word sense disambiguation preferences taken from WordNet sense frequency statistics and from the KB. In addition, we use a simple filter to discard apparently irrelevant/nonsensical assertions by discarding those that use concepts unrelated to the domain. These are defined as those with words whose preferred WordNet sense falls under one of a small number of hand-selected "non-biological" synsets (namely `human_activity#n#1`, `mental_object#n#1`, `artifact#n#1`, `instrumentation#n#1`, `psychological_feature#n#1`, `device#n#1`). One might also be able to use a KB-guided approach to disambiguation similar to that described for known facts, by (for example) looking for generalizations of (interpretations of) new facts in the knowledge base. This is a direction for future exploration.

4 Algorithm and Implementation

4.1 Topics and Participants

For now, we assume that all knowledge in the KB can be represented by "forall...exists..." statements, i.e., statements of the form:

$$\forall x \text{ isa}(x,C) \rightarrow \exists y_1..y_n p_1(v_1,v_2), \dots, p_q(v_r,v_s) \quad [1]$$

(We will later discuss how this assumption can be relaxed). p_i are predicates in the KB's ontology and each v_i is either a variable $v \in \{x, y_1, \dots, y_n\}$ or a symbol in the KB's ontology. We say clauses $p_i(v_j, v_k)$ are *about* concept C , and that C is the *topic* of the clauses. We also say that any instance y_i that is in some (possibly indirect) relationship to x is a *participant* in the representation of instance x of C . Thus all the y_i in [1] are participants, plus there may be additional participants implied by other axioms. For some given instance $X0$ of C , we can identify all the participants in the representation of $X0$ by forward chaining from $\text{isa}(X0,C)$ and collecting all the instances connected via some chain of predicates to $X0$. We encode this using a $\text{participant}(x, y_i)$ relation². As this computation is potentially unbounded, we use a depth bound to heuristically limit this computation.

4.2 Initial Language Processing

Assume the system has a paragraph of text about a topic, and that it has identified what that topic is³. For example, consider that the topic is *prophase* (a step in cell division), and the paragraph is the single sentence:

T: The mitotic spindle consists of hollow microtubules.

Text is parsed using a broad coverage, phrase structure parser (Harrison and Maxwell, 1986), followed by coreference resolution, producing a "syntactic" logical form, here:

LF: "mitotic-spindle"(s), "consist"(c), "hollow"(h), "microtubule"(m), subject(c,s), "of"(c,m), modifier(m,h).

² We can formalize this by adding $\text{participant}(x, y_i)$ to the conclusion in [1] for all y_i , plus an axiom that participant is transitive.

³ E.g., via some topic classification algorithm. For our experiments here, we manually declare the topic.

4.3 Semantic Interpretation

To interpret T the system then tries to answer, or partially answer, the corresponding question:

T: Is it true that the mitotic spindle consists of hollow microtubules?

Note that this question is in the context of the topic (prophase), and thus the mentioned objects are implicitly participants in prophase. To do this, the system proceeds as follows, using a deductive reasoning engine operating over the KB:

- (a) **setup:** create an instance X0 of the topic, i.e., assert $isa(X0, topic)$ in the KB, then find its participants $\{ y \mid participant(X0, y) \}$. Next, bind one of the variables in the LF to a participant that the variable's word might denote.
- (b) **query:** for each clause in the LF with at least one bound variable, iteratively query the KB to see if some interpretation of those clauses are provable i.e., already known.

In this example, for setup (a) the system first creates an instance X0 of prophase, i.e., asserts $isa(X0, Prophase)$ in the KB, then finds its participants Y0,...,Yn by querying for $participant(X0, y)$. The participants Y0, Y1, ... will be instances of objects and events known (in the KB) to be present in prophase, e.g., instances of Move, Centrosome, Elongate, Mitotic-Spindle, etc. The system then instantiates a variable in the LF, e.g., s in "mitotic-spindle"(s) with a participant that "mitotic spindle" might refer to, e.g., Y4, if Y4 is an instance of Mitotic-Spindle. The resulting LF looks:

LF: "mitotic-spindle"(Y4), "consist"(c), "hollow"(h), "microtubule"(m), subject(c, Y4), "of"(c, m), modifier(m, h).

For querying (b), the system uses the algorithm as follows (on the next page):

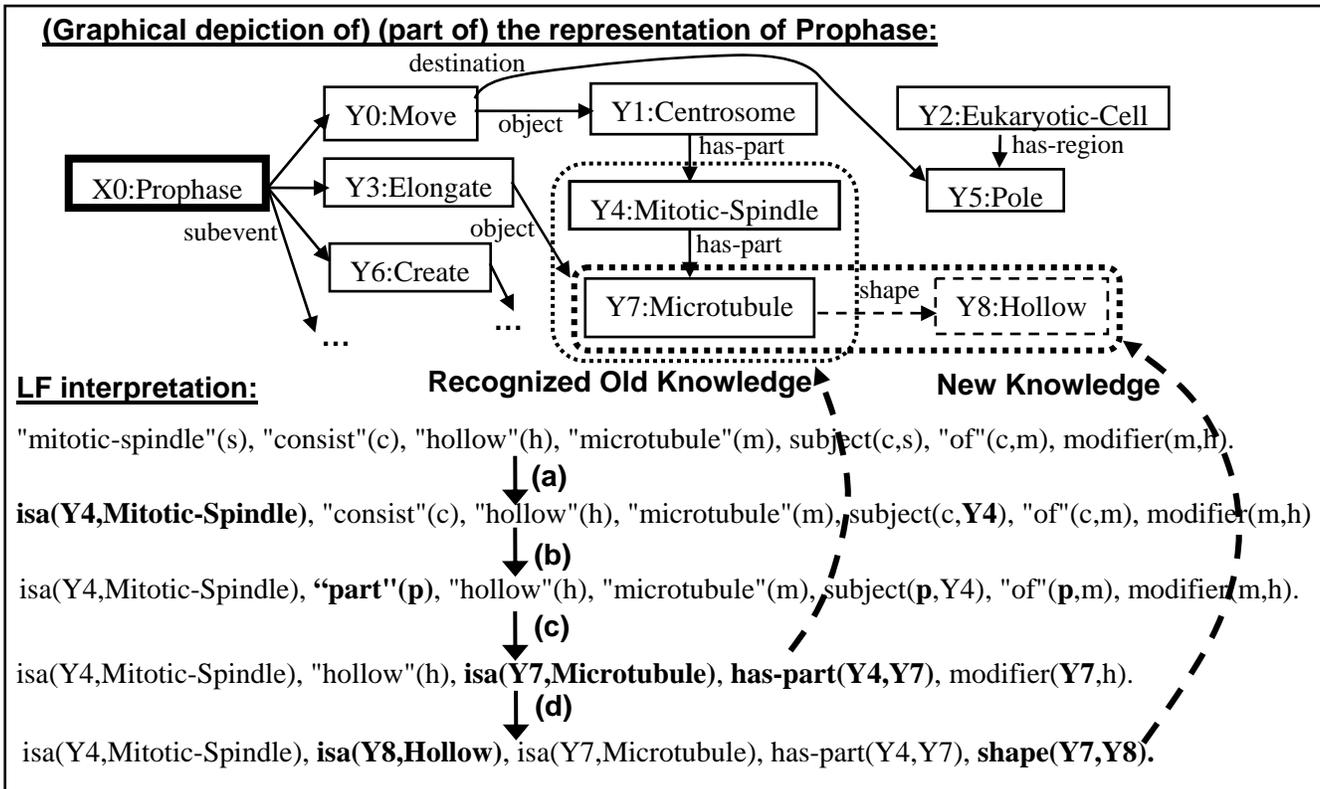


Figure 2: The path found through the search space for an interpretation of the example sentence. (a) setup (b) paraphrase substitution (c) interpretation of {subject-of(Y4, p), "part"(p), "of"(p, m)} as has-part(Y4, m), preferred as it is provable from the KB, resulting in m=Y7 (d) interpretation of new knowledge (standard WSD and SRL tools).

```

repeat
  select a clause chain  $C_u$  of "syntactic" clauses
    in the LF with at least 1 bound variable
     $C_u = \{p(x,y)\}$  or  $\{w(x)\}$  or
       $\{p_1(x,z), w(z), p_2(z,y)\}$ 
  select some interpretation  $C$  of  $C_u$  where:
     $C$  is a possible interpretation of  $C_u$ 
    or  $C'_u$  is a possible paraphrase for  $C_u$  and
       $C$  is a possible interpretation of  $C'_u$ 
  try prove  $C[\text{bindings}] \rightarrow \text{new-bindings}$ 
  If success:
    replace  $C_u$  with  $C$ 
    add new-bindings to bindings
until
  as many clauses proved as possible

```

where:

- A *syntactic clause* is a clause whose predicate is a word or syntactic role (subject, object, modifier, etc.) All clauses in the initial LF are syntactic clauses.
- A *clause chain* is a set of "syntactic" clauses in the LF of the form $\{p(x,y)\}$ or $\{w(x)\}$ or $\{p_1(x,z), w(z), p_2(z,y)\}$, where p_i, w are words or syntactic roles (subject, modifier, etc.).
- A *possible paraphrase* is a possible substitution of one syntactic clause chain with another, listed in the DIRT paraphrase database.
- A *possible interpretation* of the singleton syntactic clause chain $\{w(x)\}$ is $\text{isa}(x, \text{class})$, where *class* is a possible sense of word w .
- A *possible interpretation* of a syntactic clause chain $\{p(x,y)\}$ or $\{p_1(x,z), w(z), p_2(z,y)\}$ is $r(x,y)$, where r is a semantic relation corresponding to syntactic relation p (e.g., "in" $(x,y) \rightarrow \text{is-inside}(x,y)$) or word w (e.g., $\{\text{subject-of}(e,h), \text{"have"}(h), \text{"of"}(h,n)\} \rightarrow \text{has-part}(e,n)$).

Possible word-to-class and word-to-predicate mappings are specified in the KB.

As there are several points of non-determinism in the algorithm, including the setup (e.g., which clauses to select, which interpretation to explore), it is a search process. Our current implementation uses most-instantiated-first query ordering plus breadth-first search, although other implementations could traverse the space in other ways.

Figure 2 illustrates this procedure for the example sentence. The procedure iteratively replaces syntactic clauses with semantic clauses that corre-

spond to an interpretation that is provable from the KB. If all the clauses are proved, then the original text T is redundant; there exists an interpretation under which it can be proved from the KB, and we assume under the benevolent user assumption that this is the interpretation that the user intended.

If some syntactic clauses remain unproved, then they correspond to new knowledge, and a standard NLP pipeline is then used to interpret them. In this example (Figure 2), the "hollow" modifier to the microtubule $Y7$ was unproved, and was subsequently interpreted by the NLP pipeline as the shape of the microtubule. This new clause is converted into a (potential) addition to the KB by identifying an axiom that concluded one of the known connected facts (here, $\text{has-part}(Y4, Y7)$), and then proposing the new clause as an additional conclusion of that axiom. If there are no connected clauses, it is instead proposed as a new axiom about prophase. The user can verify/reject that proposal as he/she desires.

5 Illustration

An illustration of the system's typical processing of a paragraph is shown in Figure 3. As in Figure 1, normal font shows facts recognized as already known, and bold shows new knowledge. Again note that the output facts are not a simple recitation of the input, but have been interpreted with respect to the KB. For example, in (e), Create is the preferred interpretation of "form", and in (d), has-part is the preferred interpretation of "consisting of", as these result in the interpretation being provable from the KB. Also note that new knowledge is anchored to old, e.g., in (d), proteins are posited as an additional part of the mitotic spindle participant of prophase.

There are several errors and meaningless statements in the output also. For example, "something signals" is not particularly helpful, and "the chromosome moves" is, although biologically correct, a misinterpretation of the original English "the chromosome is seen as...", with Move being a possible interpretation of "see" (as in "I'll see you to the door"). In addition some sentences were mis-parsed or unparsed, and some interpretations were discarded as they involved non-biological concepts. Many representational issues have been skirted also, discussed shortly.

Topic: prophase

Input Paragraph:⁴

During prophase, chromosomes become visible, the nucleolus disappears, the mitotic spindle forms, and the nuclear envelope disappears. Chromosomes become more coiled and can be viewed under a light microscope. Each duplicated chromosome is seen as a pair of sister chromatids joined by the duplicated but unseparated centromere. The nucleolus disappears during prophase. In the cytoplasm, the mitotic spindle, consisting of microtubules and other proteins, forms between the two pairs of centrioles as they migrate to opposite poles of the cell. The nuclear envelope disappears at the end of prophase. This signals the beginning of the substage called prometaphase.

Output Axioms: (expressed in English)

In all prophase events:

- d. The chromosome moves.
- e. The chromatids are attached **by the centromere**.
- f. **The nucleolus disappears during the prophase**.
- g. The mitotic spindle has parts the microtubule **and the protein**.
- h. The mitotic spindle is created **between the centrioles in the cytoplasm**.
- i. The centrioles move **to the poles**.
- j. **The nuclear envelope disappears at the end**.
- k. **Something signals**.

Figure 3: Illustration of the System's Behavior

6 Preliminary Evaluation

To make a preliminary assessment of how much useful information the system is producing, we conducted a small study. 10 paragraphs about prophase (from different Web sources) were run through the system (110 sentences in total). The system extracted 114 statements of which 23 (20%) were interpreted as fully known (i.e., already in the KB), 27 (24%) as partially new knowledge, and 64 (56%) as completely new knowledge. The extracted statements were then scored by a biologist as one of:

c = correct; useful knowledge that should be in the KB

⁴ From http://www.phschool.com/science/biology_place/bio-coach/mitosisig/prophase.html

q = questionable; not useful knowledge (meaningless, overly general, vague)

i = incorrect

The results are shown in Table 1.

	Statements that are:		
	Fully known	Mixture of known & new	Fully new
Correct	22	19	25
Questionable	1	8	38
Incorrect	0	0	1

Table 1: Correctness of axioms proposed by the system.

For the statements that mix old and new knowledge, 70% were judged correct, and for completely new statements, 39% were judged correct.⁵ This suggests the system is at least producing some useful suggestions, and for the statements that mix old and new knowledge, has identified the connection points in the KB for the new facts. Although this level of accuracy is too low for automation, it suggests the system might be a useful tool for helping a knowledge engineer check that he/she has fully encoded the contents of a passage when building the KB, and performing those approved additions automatically.

7 Discussion and Conclusion

We have described a method for reading "at the fringe" of what is known, leveraging existing knowledge to help in the reading process by treating text interpretation as partial question answering. This approach is appropriate for situations in which a reasonable proportion of the text's content is already known to (represented in) the KB. Our evaluation suggests that the approach has merit, at least as an interactive knowledge acquisition tool.

As we have suggested, existing knowledge can help guide and anchor interpretation, but to what extent might it stifle the system from learning genuinely new knowledge? At present, our system is unable to extend its own ontology (it can only learn axioms expressed using its existing ontology), and thus will skim over unrecognized words

⁵ For the 1 statement already fully known but judged as questionable, the score appears to be due to poor rendering in English, the axiom being rendered as "The membrane break down." rather than "The membrane breaks down."

even if those words reflect something new (with respect to the KB) and important about the domain. The system is thus biased towards texts about concepts that it has at least heard of before (even if it knows little about them), expecting small, incremental additions of knowledge, rather than working hard to untangle information about completely novel topics. It can learn about concepts it has already heard of, but not, at present, learn new concepts. While it would be simple to modify the system to treat any new word as a new concept, this may potentially overwhelm the system, and so such extensions would need to be made carefully. This is an area for future work.

How large must the starting KB be? Although it can be missing (possibly many) axioms, we implicitly assume that at least the basic ontology and the mapping from words to concepts is reasonably complete (for the types of texts being considered), i.e., there is a good "skeleton" KB to add axioms to. Thus, methodologically, a first step for using our system would be to create the initial ontology and lexical mappings, e.g., via corpus analysis or using an ontology learning tool (Gomez-Perez and Manzano-Macho, 2003). Beyond that, the more axioms the starting KB has the better, as each axiom can potentially guide the interpretation of a new sentence. In the limiting case, where there are no axioms (only the ontology and lexical mappings), our system's behavior reverts to that of a normal, pipelined NLP interpreter (with the normal associated problems).

This work is still preliminary, and there are numerous other issues and limitations that need to be addressed also. Three notable issues are as follows:

- **Syntactic ambiguity:** While we defer WSD and SRL commitment, our system is eagerly committing to a single parse during initial language processing, and that parse may be wrong. An obvious extension is to similarly defer some syntactic commitments until semantic interpretation, for example using an underspecified or packed logical form (e.g., Bobrow et al, 2005) or exploring alternative parses.
- **Interpretation of new knowledge:** While our approach leverages the KB to interpret statements about known facts, and thus help find the anchor points for new facts, those state-

ments of new facts are still interpreted using a traditional pipelined approach, with all its associated brittlenesses (as evidenced in the last column in Table 1). Creative ways for using the KB to similarly help guide new fact interpretation are needed, for example searching the KB for generalizations or variants of those facts, and then preferring the interpretations they suggest.

- **Representational adequacy:** Our work so far has assumed a simple, deductive representational framework of individual objects and events, and correspondingly assumes the same individuality in language. However the world, and language used to describe it, often goes beyond this to include a myriad of representationally difficult phenomena (sets, pairs, aggregates, conditionality, plausibility, constraints, etc.). Our system largely skips over such aspects, as it is unable to represent them.

Despite these limitations, the picture of text interpretation as partial question-answering appears to be a useful one, as it suggests a means by which language and knowledge can be connected. We are optimistic that it can be further developed and improved for better machine reading in the future.

Acknowledgements

We are grateful to Vulcan Inc. who partially supported this work through Project Halo.

References

- Agirre, E., Marquez, L., Wicentowski, R., Eds., 2007. *Proceedings of the 4th International Workshop on Semantic Evaluations (SemEval)*, ACL, Prague, Czech Republic.
- Bentivogli, L., Dagan, I., Dang, Hoa, Giampiccolo, D., Magnini, B. 2009. The Fifth PASCAL Recognizing Textual Entailment Challenge. In *Proc Text Analysis Conference (TAC'09)*.
- Bobrow, D. G., Condoravdi, Crouch, R. Kaplan, R. Karttunen, L., King, L.T.H., de Paiva, V., Zaenen, A. 2005. A Basic Logic for Textual Inference. In *Proceedings of the AAAI Workshop on Inference for Textual Question Answering*, Pittsburgh, PA.
- Carlson, A., Betteridge, J., Hruschka, E., Mitchell, T. 2009. Coupling Semi-Supervised Learning of Cate-

- gories and Relations. *Proceedings of the NAACL HLT 2009 Workshop on Semi-supervised Learning for Natural Language Processing*.
- Clark, P., Harrison, P. 2010. Exploiting Paraphrases and Deferred Sense Commitment to Interpret Questions more Reliably. (Submitted).
- Dahlgren, K., Lord, C., Wada, H., McDowell, J., Sabler, E. 1991. ITP: Description of the intertext system as used for MUC-3. In *Proc 3rd Message Understanding Conference (MUC-3)*, pp163-170.
- DeJong, G. 1979. Prediction and Substantiation: Two processes that comprise understanding. *IJCAI'79*, pp217-222.
- Dras, M., Yamamoto, K. (Eds). 2005. *Proc 3rd International Workshop of Paraphrasing*. South Korea.
- Erk, K., Pado, S. 2006. Shalmaneser – a flexible toolbox for semantic role assignment. *Proc LREC 2006*, Genoa, Italy.
- Feigenbaum, E. 2003. Some Challenges and Grand Challenges for Computational Intelligence. *Journal of ACM*, 50 (1), pp 32-40.
- Gomez-Perez, A., Manzano-Macho, D. 2003. "A Survey of Ontology Learning Methods and Techniques", Technical Report, Univ Politecnica de Madrid (OntoWeb Deliverable 1.5, <http://www.sti-innsbruck.at/fileadmin/documents/deliverables/Ontoweb/D1.5.pdf>)
- Gunning, D., Greaves, M., Chaudhri, V. et al. 2010. *Project Halo Update – Progress Towards Digital Aristotle* (Submitted).
- Harrison, P., Maxwell, M. 1986. A New Implementation of GPSG. In *Proc. 6th Canadian Conf on AI*.
- Hobb, J., Stickel, M., Appelt, D., Martin, P. 1993. Interpretation as Abduction. In *Artificial Intelligence* 63 (1-2), pp 69-142.
- Lin, D. and Pantel, P. 2001a. Discovery of Inference Rules for Question Answering. *Natural Language Engineering* 7 (4) pp 343-360.
- Lin, D. and Pantel, P. 2001b. Discovery of Inference Rules in Text. *Proc ACM SIGKDD Conf on Knowledge Discovery and Data Mining*.
- Mulkar, R., Hobbs, J. Hovy, E. 2007. Learning to Read Syntactically Complex Biology Texts. In *Proc 8th International Symposium on Logical Formalizations of Commonsense Reasoning (AAAI Spring Symposium)*.
- Pinkal, M. 1999. On Semantic Underspecification. In Bunt, H./Muskens, R. (Eds.). *Proceedings of the 2nd International Workshop on Computational Linguistics (IWCS 2)*.
- Riloff, E. 1996. Automatically Generating Extraction Patterns from Untagged Text. *Proc AAAI'96*.
- Sekine, S. 2006. On-Demand Information Extraction. *Proc. COLING-ACL*.
- Toutanova, K., Haghighi, A., Manning, C. 2008. A Global Joint Model for Semantic Role Labeling. *Computational Linguistics*, 34 (2).
- Zadrozny, W., Jensen, K. 1991. Semantics of Paragraphs. in *Computational Linguistics* 17 (2) pp171-209.