

Open Knowledge Extraction through Compositional Language Processing

Benjamin Van Durme

University of Rochester

Rochester, NY 14627

USA

vandurme@cs.rochester.edu

Lenhart Schubert

University of Rochester

Rochester, NY 14627

USA

schubert@cs.rochester.edu

Abstract

We present results for a system designed to perform *Open Knowledge Extraction*, based on a tradition of compositional language processing, as applied to a large collection of text derived from the Web. Evaluation through manual assessment shows that well-formed propositions of reasonable quality, representing general world knowledge, given in a logical form potentially useable for inference, may be extracted in high volume from arbitrary input sentences. We compare these results with those obtained in recent work on Open *Information* Extraction, indicating with some examples the quite different kinds of output obtained by the two approaches. Finally, we observe that portions of the extracted knowledge are comparable to results of recent work on *class attribute* extraction.

1 Introduction

Several early studies in large-scale text processing, such as (Liakata and Pulman, 2002), (Gildea and Palmer, 2002), and (Schubert, 2002) showed that having access to a sentence’s syntax enabled credible, automated semantic analysis. These studies suggest that the use of increasingly sophisticated linguistic analysis tools could enable an explosion in available symbolic knowledge. Nonetheless, much of the subsequent work in extraction has remained averse to the use of the linguistic deep structure of text; this decision is typically justified by a desire to keep the extraction system as computationally lightweight as possible.

The acquisition of background knowledge is not an activity that needs to occur online; we argue that as long as the extractor will finish in a *reasonable* period of time, the speed of such a system is an issue of secondary importance. Accuracy and usefulness of knowledge

should be of paramount concern, especially as the increase in available computational power makes such “heavy” processing less of an issue.

The system explored in this paper is designed for *Open Knowledge Extraction*: the conversion of arbitrary input sentences into general world knowledge represented in a logical form possibly usable for inference. Results show the feasibility of extraction via the use of sophisticated natural language processing as applied to web texts.

2 Previous Work

Given that the concern here is with *open* knowledge extraction, the myriad projects that target a few prespecified types of relations occurring in a large corpus are set aside.

Among early efforts, one might count work on deriving selectional preferences (e.g., (Zernik, 1992; Resnik, 1993; Clark and Weir, 1999)) or partial predicate-argument structure (e.g., (Abney, 1996)) as steps in the direction of open knowledge extraction, though typically few of the tuples obtained (often a type of subject plus a verb, or a verb plus a type of object) can be interpreted as complete items of world knowledge. Another somewhat relevant line of research was initiated by Zelle and Mooney (1996), concerned with learning to map NL database queries into formal db queries (a kind of semantic interpretation). This was pursued further, for instance, by Zettlemoyer and Collins (2005) and Wong and Mooney (2007), aimed at learning log-linear models, or (in the latter case) synchronous CF grammars augmented with lambda operators, for mapping English queries to db queries. However, this approach requires annotation of texts with logical forms, and extending this approach to general texts would seemingly require a massive corpus of hand-annotated text – and the logical forms would have to cover far more phenomena than are found in db queries (e.g., attitudes, generalized quantifiers, etc.).

Another line of at least tangentially relevant

work is that on semantic role labelling. One early example was MindNet (Richardson et al., 1998), which was based on collecting 24 semantic role relations from MRDs such as the *American Heritage Dictionary*. More recent representative efforts includes that of Gildea and Jurafsky (2002), Gildea and Palmer (2002), and Punyakanok et al. (2007). The relevance of this work comes from the fact that identifying the arguments of the verbs in a sentence is a first step towards forming predications, and these may in many cases correspond to items of world knowledge.

Liakata and Pulman (2002) built a system for recovering Davidsonian predicate-argument structures from the Penn Treebank through the application of a small set of syntactic templates targeting head nodes of verb arguments. The authors illustrate their results for the sentence “*Apple II owners, for example, had to use their television sets as screens and stored data on audiocassettes*” (along with the Treebank annotations); they obtain the following QLF, where verb stems serve as predicates, and arguments are represented by the head words of the source phrases:

```
have(e1,owner, (use(e3,owner,set),
               and as(e3,screen)),
and
(store(e2,owner,datum),
 and on(e2,audiocassette)))
```

For a test set of 100 Treebank sentences, the authors report recall figures for various aspects of such QLFs ranging from 87% to 96%. While a QLF like the one above cannot in itself be regarded as world knowledge, one can readily imagine postprocessing steps that could in many cases obtain credible propositions from such QLFs. How accurate the results would be with machine-parsed sentences is at this point unknown.

In the same year, Schubert (2002) described a project aimed directly at the extraction of general world knowledge from Treebank text, and Schubert and Tong (2003) provided the results of hand-assessment of the resulting propositions. The Brown corpus yielded about 117,000 distinct simple propositions (somewhat more than 2 per sentence, of variable quality). Like Liakata and Pulman’s approach the method relied on the computation of unscoped logical forms from Treebank trees, but it abstracted propositional information along the way, typically discarding modifiers at deeper levels from LFs at higher levels, and also replacing NPs (including named entities) by their types as far as possible. Judges found about 2/3 of the output propositions (when automatically verbalized in English) acceptable as general

claims about the world. The next section provides more detail on the extraction system, called KNEXT, employed in this work.

Clark et al. (2003), citing the 2002 work of Schubert, report undertaking a similar extraction effort for the 2003 Reuters corpus, based on parses produced by the Boeing parser, (see (Holmback et al., 2000)), and obtained 1.1 million subject-verb-object fragments. Their goal was eventually to employ such tuples as common-sense expectations to guide the interpretation of text and the retrieval of possibly relevant knowledge in question-answering. This goal, unlike the goal of inferential use of extracted knowledge, does not necessarily require the extracted information to be in the form of logical propositions. Still, since many of their tuples were in a form that could be quite directly converted into propositional forms similar to those of Schubert, their work indicated the potential for scalability in parser-based approaches to information extraction or knowledge extraction.

A recent project aimed at large-scale, open extraction of tuples of text fragments representing verbal predicates and their arguments is TextRunner (Banko et al., 2007). This system does part-of-speech tagging of a corpus, identifies noun phrases with a noun phrase chunker, and then uses tuples of nearby noun phrases within sentences to form apparent relations, using intervening material to represent the relation. Apparent modifiers such as prepositional phrases after a noun or adverbs are dropped. Every candidate relational tuple is classified as trustworthy (or not) by a Bayesian classifier, using such features as parts of speech, number of relevant words between the noun phrases, etc. The Bayesian classifier is obtained through training on a parsed corpus, where a set of heuristic rules determine the trustworthiness of apparent relations between noun phrases in that corpus. As a preview of an example we will discuss later, here are two relational tuples in the format extracted by TextRunner:

```
(the people) use (force),
(the people) use (force) to impose (a government).
```

(Boldface indicates items recognized as head nouns.) No attempt is made to convert text fragments such as “*the people*” or “*use – to impose*” into logically formal terms or predicates. Thus much like semantic role-labelling systems, TextRunner is an *information* extraction system, under the terminology used here; however, it comes

closer to knowledge extraction than the former, in that it often strips away much of the modifying information of complex terms (e.g., leaving just a head noun phrase).

2.1 KNEXT

KNEXT (Schubert, 2002) was originally designed for application to collections of manually annotated parse trees, such as the Brown corpus. In order to extract knowledge from larger text collections, the system has been extended for processing arbitrary text through the use of third-party parsers. In addition, numerous improvements have been made to the semantic interpretation rules, the filtering techniques, and other components of the system. The extraction procedure is as follows:

1. Parse each sentence using a Treebank-trained parser (Collins, 1997; Charniak, 1999).
2. Preprocess the parse tree, for better interpretability (e.g., distinguish different types of SBAR phrases and different types of PPs, identify temporal phrases, etc.).
3. Apply a set of 80 interpretive rules for computing unscoped logical forms (ULFs) of the sentence and all lower-level constituents in a bottom-up sweep; at the same time, *abstract* and collect phrasal logical forms that promise to yield stand-alone propositions (e.g., ULFs of clauses and of pre- or post-modified nominals are prime candidates). The ULFs are rendered in Episodic Logic (e.g., (Schubert and Hwang, 2000)), a highly expressive representation allowing for generalized quantifiers, predicate modifiers, predicate and sentence reification operators, and other devices found in NL. The abstraction process drops modifiers present in lower-level ULFs (e.g., adjectival premodifiers of nominal predicates) in constructing higher-level ULFs (e.g., for clauses). In addition, named entities are generalized as far as possible using several gazetteers (e.g., for male and female given names, US states, world cities, actors, etc.) and some morphological processing.
4. Construct complete sentential ULFs from the phrasal ULFs collected in the previous step; here some filtering is performed to exclude vacuous or ill-formed results.
5. Render the propositions from the previous step in (approximate) English; again significant heuristic filtering is done here.

As an example of KNEXT output, the sentence

Cock fights, however, are still legal in six of the United States, perhaps because we still eat chicken regularly,

but no-longer dogs

yields a pair of propositions expressed logically as
(:i (:f k (:f nm cock[n] (:f plur fight[n]))) legal[a]),
(:i (:q det (:f plur person)) eat[v] (:f k chicken[n]))
and these are automatically rendered in approximate English as

COCK FIGHTS CAN BE LEGAL.

PERSON -S MAY EAT CHICKEN.

Additional examples will be seen later, with logical forms translated into a more readable, conventional format.

One larger collection we have processed since the 2002-3 work on Treebank corpora is the British National Corpus (BNC), consisting of 200 million words of mixed-genre text passages. (The above example is taken from there.) The quality of resulting propositions has been assessed by the hand-judging methodology of Schubert and Tong (2003), yielding positive judgements almost as frequently as for the Brown Treebank corpus. The next section, concerned with the web corpus collected and used by Banko et al. (2007), contains a fuller description of the judging method. The BNC-based KB, containing 6,205,877 extracted propositions, is publicly searchable via a recently developed online knowledge browser.¹

3 Experiments

The experiments reported here were aimed at a comparative assessment of linguistically based knowledge extraction (by KNEXT), and pattern-based information extraction (by TextRunner, and by another system, aimed at class attribute discovery).

Dataset Experiments were based on sampling 1% of the sentences from each document contained within a corpus of 11,684,774 web pages harvested from 1,354,123 unique top level domains. The top ten contributing domains can be seen in table 1, which together make up 31.5% of the documents in the collection. There were 310,463,012 sentences in all, the sample containing 3,000,736. Of these, 1,373 were longer than a preset limit of 100 tokens, and were discarded.² Sentences containing individual tokens of length greater than 500 characters were similarly removed.³

¹<http://www.cs.rochester.edu/u/vandurme/epik>

²Typically enumerations, e.g., *There have been 29 MET deployments in the city of Florida since the inception of the program : three in Ft. Pierce , Collier County , Opa Locka , ...*

³For example, *Kellnull phenotypes can occur through splice site and splice-site / frameshift mutations301,302*

domain	freq.	%	type	domain	freq.	%	type
en.wikipedia.org	1,616,279	13.8	ref	findarticles.com	56,173	0.5	misc
www.answers.com	1,318,097	11.3	ref	www.geocities.com	52,262	0.4	misc
www.amazon.com	257,868	2.2	shop	www.city-data.com	50,891	0.4	ref
www.imdb.com	182,087	1.6	ent	www.tv.com	41,699	0.4	ent
www.britannica.com	59,269	0.5	ref	www.cduniverse.com	40,859	0.3	shop

Table 1: Top ten most frequent domains and their relative percent contribution to the corpus (by document count), with *type* being one of: reference material (ref), online shopping (shop), entertainment facts or news (ent), and miscellaneous (misc).

1.	A REASONABLE GENERAL CLAIM e.g., A grand-jury may say a proposition
2.	TRUE BUT TOO SPECIFIC TO BE USEFUL e.g., Bunker walls may be decorated with seashells
3.	TRUE BUT TOO GENERAL TO BE USEFUL e.g., A person can be nearest an entity
4.	SEEMS FALSE e.g., A square can be round
5.	SOMETHING IS OBVIOUSLY MISSING e.g., A person may ask
6.	HARD TO JUDGE e.g., Supervision can be with a company

Figure 1: Instructions for categorical judging.

The corpus was a superset of that used by Banko et al. (2007); for each sentence in the collection there exist zero or more *tuples* as extracted by the TextRunner system.

Note that while websites such as Wikipedia.org contain large quantities of structured information stored in lists and tables, the focus here is entirely on natural language sentences. In addition, as the extraction methods discussed in this paper do not make use of intersentential features, the lack of sentence to sentence coherence resulting from random sampling had no effect on the results.

Extraction Sentences were processed using the syntactic parser of Charniak (1999). From the resultant trees, KNEXT extracted 7,406,371 *propositions*, giving a *raw* average of 2.47 per sentence. Of these, 4,151,779 were unique, so that the average extraction frequency per sentence is 1.78 unique propositions. Post-processing left 3,975,197 items, giving a per sentence expectation of 1.32 unique, filtered propositions. Selected examples regarding knowledge about people appear in table 2.

For the same sample, TextRunner extracted 6,053,983 *tuples*, leading to a *raw* average of 2.02 tuples per sentence. As described by its designers, TextRunner is an *information* extraction system; while KNEXT might seem to “win” in raw extraction volume, these numbers are not in fact directly comparable (see section on *Comparison*).

Evaluation Extraction quality was determined 450039003[...]3000 premature stop codons and missense mutations.

Category	% Selected	Kappa	% Selected	Kappa
1	49%	0.4017	50%	0.2822
1, 2, or 3	54%	0.4766	60%	0.3360
		judges	judges w/	volunteers

Table 3: Percent propositions labeled under the given category(s), paired with Fleiss’ Kappa scores. Results are reported both for the authors (judges one and two), along with two volunteers.

THE STATEMENT ABOVE IS A REASONABLY CLEAR, ENTIRELY PLAUSIBLE GENERAL CLAIM AND SEEMS NEITHER TOO SPECIFIC NOR TOO GENERAL OR VAGUE TO BE USEFUL:
1. I agree.
2. I lean towards agreement.
3. I’m not sure.
4. I lean towards disagreement.
5. I disagree.

Figure 2: Instructions for scaled judging.

through manual assessment of verbalized propositions drawn randomly from the results. Initial evaluation was done using the method proposed in Schubert and Tong (2003), in which judges were asked to label propositions according to their category of acceptability; abbreviated instructions may be seen in figure 1.⁴ Under this framework, category one corresponds to a strict assessment of acceptability, while an assignment to any of the categories between one and three may be interpreted as a weaker level of acceptance. As seen in table 3, average acceptability was judged to be roughly 50 to 60%, with associated Kappa scores signalling fair (0.28) to moderate (0.48) agreement.

Judgement categories at this level of specificity are useful both for system analysis at the development stage, as well as for training judges to recognize the disparate ways in which a proposition may not be acceptable. However, due to the rates of agreement observed, evaluation moved to the use of a five point sliding scale (figure 2). This scale allows for only a single axis of comparison, thus collapsing the various ways in which a proposition may or may not be flawed into a single,

⁴Judges consisted of the authors and two volunteers, each with a background in linguistics and knowledge representation.

A PERSON ...		
MAY SING TO A GIRLFRIEND	MAY RECEIVE AN ORDER FROM A GENERAL	CAN BE WITHOUT PRETENTIONS
MAY BURN A SAWMILL	MAY CARRY IMAGES OF A WOMAN	MAY RESPOND TO A QUESTION
MAY WALK WITH A FRIEND	MAY CHAT WITH A MALE-INDIVIDUAL	MAY FEIGN A DISABILITY
MAY PICK UP A PHONE	MAY LOOK FOR A (QUALITY SHAMPOO PRODUCT)	
MAY DOWNLOAD AN ALBUM	MAY MUSH A TEAM OF (SEASONED SLED DOGS)	
MAY SING TO A GIRLFRIEND	MAY OBTAIN SOME NUMBER OF (PERCULA CLOWNFISH)	
MAY BUY FOOD	MAY KNOW STUFF	MAY PRESENT A PAPER
MAY LIKE (POP CULTURE)	MAY SHOOT A MALE-INDIVIDUAL	MAY EXPERIENCE A FEELING

Table 2: Verbalized propositions concerning the class PERSON.

general notion of acceptability.

The authors judged 480 propositions sampled randomly from amongst bins corresponding to frequency of support (i.e., the number of times a given proposition was extracted). 60 propositions were sampled from each of 8 such ranges.⁵ As seen in the first graph of figure 3, propositions that were extracted at least twice were judged to be more acceptable than those extracted only once. While this is to be expected, it is striking that as frequency of support increased further, the level of judged acceptability remained roughly the same.

4 Comparison

To highlight differences between an extraction system targeting knowledge (represented as logical statements) as compared to information (represented as segmented text fragments), the output of KNEXT is compared to that of TextRunner for two select inputs.

4.1 Basic

A defining quote from the book, “An armed society is a polite society”, is very popular with those in the United States who support the personal right to bear arms.

From this sentence TextRunner extracts the tuples:⁶

(A defining **quote**) is a (polite **society**),
(the personal **right**) to bear (**arms**).

We might manually translate this into a crude sort of logical form:

IS-A(A-DEFINING-QUOTE, POLITE-SOCIETY-”),
TO-BEAR(THE-PERSONAL-RIGHT, ARMS).

Better would be to consider only those terms classified as head, and make the assumption that each tuple argument implicitly introduces its own quantified variable:

$\exists x, y. \text{QUOTE}(x) \ \& \ \text{SOCIETY}(y) \ \& \ \text{IS-A}(x, y),$
 $\exists x, y. \text{RIGHT}(x) \ \& \ \text{ARMS}(y) \ \& \ \text{TO-BEAR}(x, y).$

⁵(0, 2⁰, 2¹, 2³, 2⁴, 2⁶, 2⁸, 2¹⁰, 2¹²), i.e., (0,1], (1,2], (2,8], ...

⁶Tuple *arguments* are enclosed in parenthesis, with the items recognized as *head* given in bold. All non-enclosed, conjoining text makes up the tuple *predicate*.

Compare this to the output of KNEXT:⁷

$\exists x. \text{SOCIETY}(x) \ \& \ \text{POLITE}(x),$
 $\exists x, y, z. \text{THING-REFERRED-TO}(x) \ \& \ \text{COUNTRY}(y)$
 $\ \& \ \text{EXEMPLAR-OF}(z, y) \ \& \ \text{IN}(x, z),$
 $\exists x. \text{RIGHT}(x) \ \& \ \text{PERSONAL}(x),$
 $\exists x, y. \text{QUOTE}(x) \ \& \ \text{BOOK}(y) \ \& \ \text{FROM}(x, y),$
 $\exists x. \text{SOCIETY}(x) \ \& \ \text{ARMED}(x),$

which is automatically verbalized as:

A SOCIETY CAN BE POLITE,
A THING-REFERRED-TO CAN BE IN
AN EXEMPLAR-OF A COUNTRY,
A RIGHT CAN BE PERSONAL,
A QUOTE CAN BE FROM A BOOK,
A SOCIETY CAN BE ARMED.

4.2 Extended Tuples

While KNEXT uniquely recognizes, e.g., adjectival modification and various types of possessive constructions, TextRunner more aggressively captures tuples of extended cardinality. For example, from the following:

James Harrington in The Commonwealth of Oceana uses the term anarchy to describe a situation where the people use force to impose a government on an economic base composed of either solitary land ownership, or land in the ownership of a few.

TextRunner extracts 19 tuples, some with three or even four arguments, thus aiming beyond the binary relations that most current systems are limited to. That so many tuples were extracted for a single sentence is explained by the fact that for most tuples containing $N > 2$ arguments, TextRunner will also output the same tuple with $N - 1$ arguments, such as:

(the **people**) use (**force**),
(the **people**) use (**force**) to impose (a **government**),
(the **people**) use (**force**) to impose (a **government**)
on (an economic **base**).

In addition, tuples may overlap, without one being a proper subset of another:

⁷For expository reasons, scoped, simplified versions of KNEXT’s ULFs are shown. More accurately propositions are viewed as weak *generic conditionals*, with a non-zero lower bound on conditional frequency, e.g., $[\exists x. \text{QUOTE}(x)] \Rightarrow_{0.1} [\exists y. \text{BOOK}(y) \ \& \ \text{FROM}(x, y)]$, where x is dynamically bound in the consequent.

(a **situation**) where (the **people**) use (**force**),
 (**force**) to impose (a **government**),
 (a **government**) on (an economic **base**) composed of
 (either solitary land **ownership**).

This overlap raises the question of how to accurately quantify system performance. When measuring average extraction quality, should samples be drawn randomly across tuples, or from originating sentences? If from tuples, then sample sets will be biased (for good or ill) towards fragments derived from complex syntactic constructions. If sentence based, the system fails to be rewarded for extracting as much from an input as possible, as it may conservatively target only those constructions most likely to be correct. With regards to volume, it is not clear whether adjuncts should each give rise to additional facts added to a final total; optimal would be the recognition of such optionality; failing this perhaps a tally may be based on unique predicate head terms?

As a point of merit according to its designers, TextRunner does not utilize a parser (though as mentioned it does part of speech tagging and noun phrase chunking). This is said to be justified in view of the known difficulties in reliably parsing open domain text as well as the additional computational costs. However, a serious consequence of ignoring syntactic structure is that incorrect bracketing across clausal boundaries becomes all too likely, as seen for instance in the following tuple:

(**James Harrington**) uses (the term **anarchy**) to describe (a **situation**) where (the **people**),

or in the earlier example where *from the book*, “*An armed society* appears to have been erroneously treated as a post-nominal modifier, intervening between the first argument and the *is-a* predicate.

KNEXT extracted the following six propositions, the first of which was automatically filtered in post-processing for being overly vague:⁸

*A MALE-INDIVIDUAL CAN BE IN A NAMED-ENTITY
 OF A NAMED-ENTITY,
 A MALE-INDIVIDUAL MAY USE A (TERM ANARCHY),
 PERSONS MAY USE FORCE,
 A BASE MAY BE COMPOSED IN SOME WAY,
 A BASE CAN BE ECONOMIC,
 A (LAND OWNERSHIP) CAN BE SOLITARY.

5 Extracting from Core Sentences

We have noted the common argument against the use of syntactic analysis when performing large-scale extraction viz. that it is too time consuming

⁸The authors judge the third, fifth and sixth propositions to be both well-formed and useful.

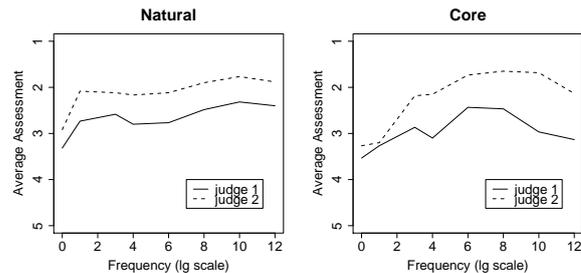


Figure 3: As a function of frequency of support, average assessment for propositions derived from *natural* and *core* sentences.

	natural	core	overlap
judge 1	3.35	3.85	2.96
judge 2	2.95	3.59	2.55

Table 4: Mean judgements (lower is better) on propositions sampled from those supported either exclusively by natural or core sentences, or those supported by both.

to be worthwhile. We are skeptical of such a view, but decided to investigate whether an argument-bracketing system such as TextRunner might be used as an extraction *preprocessor* to limit what needed to be parsed.

For each TextRunner tuple extracted from the sampled corpus, *core* sentences were constructed from the predicate and noun phrase arguments,⁹ which were then used as input to KNEXT for extraction.

From 6,053,981 tuples came an equivalent number of core sentences. Note that since TextRunner tuples may overlap, use of these reconstructed sentences may lead to skewed propositional frequencies relative to “normal” text. This bias was very much in evidence in the fact that of the 10,507,573 propositions extracted from the core sentences, only 3,787,701 remained after automatic postprocessing and elimination of duplicates. This gives a per-sentence average of 0.63, as compared to 1.32 for the original text.

While the raw number of propositions extracted for each version of the underlying data look similar, 3,975,197 (natural) vs. 3,787,701 (core), the actual overlap was less than would be expected. Just 2,163,377 propositions were extracted jointly from both natural and core sentences, representing a percent overlap of 54% and 57% respectively.

Quality was evaluated by each judge assessing 240 randomly sampled propositions for each of: those extracted exclusively from natural sen-

⁹Minor automated heuristics were used to recover, e.g., missing articles dropped during tuple construction.

tences, those extracted exclusively from core sentences, those extracted from both (table 4). Results show that propositions exclusively derived from core sentences were most likely to be judged poorly. Propositions obtained both by KNEXT alone and by KNEXT- processing of TextRunner-derived core sentences (the overlap set) were particularly likely to be judged favorably.

On the one hand, many sentential fragments ignored by TextRunner yield KNEXT propositions; on the other, TextRunner’s output may be assembled to produce sentences yielding propositions that KNEXT otherwise would have missed. Ad-hoc analysis suggests these new propositions derived with the help of TextRunner are a mix of noise stemming from bad tuples (usually a result of the aforementioned incorrect clausal bracketing), along with genuinely useful propositions coming from sentences with constructions such as appositives or conjunctive enumerations where TextRunner outguessed the syntactic parser as to the correct argument layout. Future work may consider whether (syntactic) language models can be used to help prune core sentences before being given to KNEXT.

The second graph of figure 3 differs from the first at low frequency of support. This is the result of the partially redundant tuples extracted by TextRunner for complex sentences; the core verb-argument structures are those most likely to be correctly interpreted by KNEXT, while also being those most likely to be repeated across tuples for the same sentence.

6 Class Properties

While TextRunner is perhaps the extraction system most closely related to KNEXT in terms of generality, it occurred to us that there is also significant potential overlap with work on *class attribute* extraction. Paşca and Van Durme (2007) recently described this task, going on to detail

COUNTRY	government, war, team, history, rest, coast, census, economy, <i>population</i> , independence
*DRUG	<i>side effects</i> , influence, <i>uses</i> , doses, manufacturer, efficacy, release, graduates, plasma levels, safety
*CITY	makeup, heart, center, <i>population</i> , history, side, places, name, edge, area
*PAINTER	<i>works</i> , art, brush, skill, lives, sons, friend, order quantity, muse, eye
COMPANY	windows, products, word, page, review, film, team, award, studio, director

Table 5: By frequency, the top ten attributes a class MAY HAVE. Emphasis added to entries overlapping with those reported by Paşca and Van Durme. Results for starred classes were derived without the use of prespecified lists of instances.

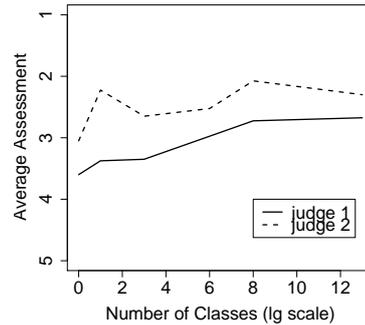


Figure 4: Mean quality of class attributes as a function of the number of classes sharing a given property.

an approach for collecting such attributes from search engine query logs. As an example, the search query “*president of Spain*” suggests that a *Country* may have a *president*.

If one were to consider attributes to correspond, at least in part, to things a class MAY HAVE, CAN BE, or MAY BE, then a subset of KNEXT’s results may be discussed in terms of this specialized task. For example, for the five classes used in those authors’ experiments, table 5 contains the top ten most frequently extracted things each class MAY HAVE, as determined by KNEXT, without any targeted filtering or adaptation to the task.

For each of these three types of attributive categories the authors judged 80 randomly drawn propositions, constrained such that half (40 for each) were supported by a single sentence, while the other half were required only to have been extracted at least twice, but potentially many hundreds or even thousands of times. As seen in table 6, the judges were strongly correlated in their assessments, where for MAY HAVE and MAY BE they were lukewarm (3.0) or better on the majority of those seen.

In a separate evaluation judges considered whether the number of classes sharing a given attribute was indicative of its acceptability. For each unique attributive proposition the class in

	1	2+	1	2+	corr.
MAY HAVE	2.8	2.35	2.5	2.28	0.68
MAY BE	3.2	2.85	2.35	2.13	0.59
CAN BE	3.78	3.58	3.28	2.75	0.76
	judge 1		judge 2		

Table 6: Mean assessed acceptability for properties that occur for only a single class (1), and more than a single class (2+). The final column gives the Pearson correlation between the two judges.

“subject” position was removed, leaving fragments such as that given in bold: **A ROBOT CAN BE SUBHUMAN**. These attribute fragments were tallied and binned by frequency,¹⁰ with 40 then sampled from each. For a given attribute selected, a single attributive proposition matching that fragment was randomly drawn. For example, having selected the attribute CAN BE FROM A US-CITY, the proposition SOME_NUMBER_OF SHERIFFS CAN BE FROM A US-CITY was drawn from the 390 classes sharing this property. As seen in figure 4, acceptability rose as a property became more common.

7 Conclusions

Work such as TextRunner (Banko et al., 2007) is pushing extraction researchers to consider larger and larger datasets. This represents significant progress towards the greater community’s goal of having access to large, expansive stores of general world knowledge.

The results presented here support the position that advances made over decades of research in parsing and semantic interpretation do have a role to play in large-scale knowledge acquisition from text. The price paid for linguistic processing is not excessive, and an advantage is the logical formality of the results, and their versatility, as indicated by the application to class attribute extraction.

Acknowledgements

We are especially grateful to Michele Banko and her colleagues for generously sharing data and results. This work was supported by NSF grants IIS-0328849 and IIS-0535105.

References

Abney, Steven. 1996. Partial parsing via finite-state cascades. In *Proc. of the ESSLLI '96 Robust Parsing Workshop*.

Banko, Michele, Michael J Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007. Open information extraction from the web. In *Proc. of IJCAI-07*, Hyderabad, India.

Charniak, Eugene. 1999. A maximum-entropy-inspired parser. Technical Report CS-99-12.

Clark, Stephen and David Weir. 1999. An iterative approach to estimating frequencies over a semantic hierarchy. In *Proc. of the Joint SIGDAT Conf. on Empirical Methods in Natural Language Processing and Very Large Corpora*.

Clark, Peter, Phil Harrison, and John Thompson. 2003. A knowledge-driven approach to text meaning processing. In *Proc. of the HLT-NAACL 2003 Workshop on Text Meaning*, pages 1–6, Morristown, NJ.

Collins, Michael. 1997. Three generative, lexicalised models for statistical parsing. In *ACL97*, pages 16–23, Madrid, Spain.

Gildea, Daniel and Daniel Jurafsky. 2002. Automatic Labeling of Semantic Roles. *Computational Linguistics*, 28(3):245–288.

Gildea, Daniel and Martha Palmer. 2002. The necessity of syntactic parsing for predicate argument recognition. In *Proc. of ACL-02*, pages 239–246, Philadelphia, PA.

Holmback, H., L. Duncan, and P. Harrison. 2000. A word sense checking application for simplified English. In *3rd Int. Workshop on Controlled Language Applications*.

Liakata, Maria and Stephen Pulman. 2002. From trees to predicate argument structures. In *Proc. of COLING 2002*, pages 563–569, Taiwan.

Paşca, Marius and Benjamin Van Durme. 2007. What you seek is what you get: Extraction of class attributes from query logs. In *Proc. of IJCAI-07*, pages 2832–7, Hyderabad, India.

Punyakanok, V., D. Roth, and W. Yih. 2007. The importance of syntactic parsing and inference in semantic role labeling. *Computational Linguistics*.

Resnik, P. 1993. Semantic classes and syntactic ambiguity. In *Proc. of ARPA Workshop on Human Language Technology*, Plainsboro, NJ.

Richardson, S. D., W. Dolan, and L. Vanderwende. 1998. MindNet: Acquiring and structuring semantic information from text. In *Proc. of COLING-ACL'98*, pages 1098 – 1102, Montreal, Quebec.

Schubert, Lenhart K. and Chung Hee Hwang. 2000. Episodic logic meets little red riding hood: A comprehensive, natural representation for language understanding. In Iwanska, L. and S.C. Shapiro, editors, *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*, pages 111–174. MIT/AAAI Press.

Schubert, Lenhart K. and Matthew H. Tong. 2003. Extracting and evaluating general world knowledge from the brown corpus. In *Proc. of the HLT/NAACL Workshop on Text Meaning*, Edmonton, Alberta, Canada.

Schubert, Lenhart K. 2002. Can we derive general world knowledge from texts? In *Proc. of HLT-02*, pages 94–97, San Diego, CA.

Wong, Yuk Wah and Raymond J. Mooney. 2007. Learning synchronous grammars for semantic parsing with lambda calculus. In *Proc. of ACL-07*, pages 960–7, Prague, Czech Republic, June.

Zelle, John M. and Raymond J. Mooney. 1996. Learning to parse database queries using inductive logic programming. In *Proc. of AAAI-96*, Portland, OR, Aug.

Zernik, Uri. 1992. Closed yesterday and closed minds: Asking the right questions of the corpus to distinguish thematic from sentential relations. In *Proc. of COLING-92*, Nantes, France, Aug. 23–28,.

Zettlemoyer, Luke and Michael Collins. 2005. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *Proc. of UAI*.

¹⁰Ranges : (0, 2⁰, 2¹, 2³, 2⁶, ∞)