Automatic Ellipsis Resolution: Recovering Covert Information from Text

Marjorie McShane & Petr Babkin
Rensselaer Polytechnic Institute
{margemc34, petr.a.babkin}@gmail.com

Abstract

Ellipsis is a linguistic process that makes certain aspects of text meaning not directly traceable to surface text elements and, therefore, inaccessible to most language processing technologies. However, detecting and resolving ellipsis is an indispensable capability for language-enabled intelligent agents. The key insight of the work presented here is that not all cases of ellipsis are equally difficult: some can be detected and resolved with high confidence even before we are able to build agents with full human-level semantic and pragmatic understanding of text. This paper describes a fully automatic, implemented and evaluated method of treating one class of ellipsis: elided scopes of modality. Our cognitively-inspired approach, which centrally leverages linguistic principles, has also been applied to overt referring expressions with equally promising results.

Introduction

The grammatical process of ellipsis – the non-expression of semantic content – increases the efficiency of natural language communication by people but introduces challenges for machine processing. For example, the final clause in sentence (1)[1] will be of little use to an intelligent agent until and unless it is decorated with metadata indicating that what the toddler *could* do is *teach illiterate women to read and write.* (In this and subsequent examples, elided categories are indicated by [e] and their sponsors – i.e., the material used to recover their meaning – are italicized.[2])

(1) Aid workers in war-ravaged Kabul were stunned when a toddler from a poor family offered to *teach illiterate women to read and write* – and then promptly proved he could [e].

This paper reports an implemented and evaluated method of automatically resolving elided scopes of modality, which expands upon the preliminary work reported in McShane et al. 2012b. The work is novel because ellipsis has been left largely untreated by natural language processing (NLP) systems (Spenader and Hendriks 2005), being resistant to the currently prevalent methodological preference for supervised machine learning. The closest point of comparison is Hardt's (1997) VP ellipsis system; however, whereas that system requires a perfect (manually corrected) syntactic parse as input, ours does not rely on any manually provided prerequisites. Other comparisons require the leap from elided to overt referring expressions; Lee et al. 2013 provides a nice overview of the current state of the art in resolving those.

The key point about ellipsis that has been overlooked is that while resolving some instances does require sophisticated semantic and pragmatic reasoning, not all cases are so difficult. Leveraging linguistic insights, our system automatically detects which instances of ellipsis it can resolve with high confidence and treats only those, supplementing each resolution decision with a corpus-attested measure of confidence that will contribute to an intelligent agent's evaluation of whether it has understood the meaning of a language input sufficiently to act in response to it.

The work is being carried out within the OntoSem language processing system of the OntoAgent cognitive architecture (McShane and Nirenburg 2012). OntoAgents are cognitively simulated, language-endowed intelligent agents that are being configured to collaborate with people in task-oriented applications, such as the prototype Maryland Virtual Patient physician training system (Nirenburg, McShane and Beale 2008). OntoAgents carry out deep semantic and pragmatic language analysis, yielding ontologically grounded text meaning representations that populate agent memory and subsequently support agent reasoning (McShane et al. 2012a; McShane, Nirenburg and Beale, Submitted).

Research and development in the OntoAgent paradigm seeks to balance the socio-economic necessity of near-term utility with the scientific goal of fundamentally solving

[1] Most examples are from the Gigaword corpus (Graff and Cieri 2003), which we used for system testing and evaluation.

[2] The inflectional form of the head of the sponsor is irrelevant since, in reality, an elided *meaning* is recovered based on the *meaning* of the sponsor. For discussion of the term *sponsor,* see Byron 2004, p. 16.

problems of artificial intelligence. In this spirit, we simultaneously work toward preparing OntoAgents to fully interpret every text input, on the one hand, and to leverage well-selected "lightweight" strategies to meet near-term goals, on the other. Our work on ellipsis detection and resolution is a good example of how those strategies can interact over time. First we configure a lightweight – though linguistically and cognitively inspired – strategy that offers useful, near-term results, albeit with incomplete coverage; then we incorporate more knowledge-intensive analysis methods, as needed, to expand coverage.

Linguistic Background

"Modality" refers to a class of meanings that can scope over a proposition (cf. recent work on sentiment analysis, e.g., Pang and Lee 2008). A given proposition can be scoped over by many different types of modal meanings, as shown by the italicized strings in (2). In addition, multiple modalities can scope over a proposition, as shown by (3).

(2) The US *has <has not, might have, cannot have, should not, might, wants to, does not want to, seems to have, could not have, is believed to have, failed to*, etc.> sign(ed) the treaty.

(3) The US *might* *have to* sign the treaty.

The microtheory of modality in OntoSem distinguishes ten types of modality: epistemic, belief, obligative, permissive, potential, evaluative, intentional, epiteuctic, effort, and volitive (Nirenburg and Raskin 2004). Values of each modality can be realized in language by various words and phrases: e.g., **volitive**: *want to, not want to*; **permissive**: *may, may not*; **epiteuctic**: *succeed in, fail to*. Each modal meaning has a scope that represents the event(s) the meaning applies to. The scope can be overtly specified or elided; if elided, it must have a *sponsor* that allows the reader/hearer to reconstruct the unexpressed meaning.

Although sponsors for elided referring expressions – like sponsors for any referring expression – are most accurately understood as semantic entities rather than text strings, the latter can serve as a proxy which is, on the one hand, sufficient for some applications and, on the other hand, a strong first step toward the more involved analysis needed to arrive at full semantic resolution. For example, in (4) the elided category refers to a different instance of *sending* a different set of *children* abroad than is indicated in the sponsor clause – a phenomenon known as *sloppy coreference* (Fiengo and May 1994).

(4) Better-off parents could *send their children abroad for English education* but poorer families could not [e].

In the system implementation reported here, the ellipsis treatment is deemed correct if the system (a) correctly detects the ellipsis and (b) selects the leftmost verbal element of the sponsor, which might be a the head of the sponsoring proposition (like *send* in (4)), or a modal or non-modal matrix verb scoping over the proposition. The notion of "leftmost verbal element" of the sponsor requires further comment.

When the sponsor-containing clause includes one or more matrix verbs, they can either be included in or excluded from the sponsor: e.g., in (4) the modal element *could* must be excluded from the sponsor. For both semantic and functional reasons, matrix verbs can be usefully divided into two subclasses: modal (*can, could, expect to,* etc.) and non-modal (*agree, ask, avoid, beg, challenge, choose,* etc.), which are treated differently by the system.

When *modal* matrix verbs appear in the sponsor clause, the decision of whether to include them in, or exclude them from, the sponsor is guided by a half dozen modality correlation rules of the following type:

- If the sponsor clause contains *effort* or *volitive* modality and the ellipsis-licensing modality is *epiteuctic* or *epistemic*, the sponsor should exclude the sponsor-clause modality: John wanted to *ski* and did [e].
- If the sponsor clause contains any number of modalities with the outer one being *epistemic* (indicating negation) and the ellipsis-licensing modality being *epistemic* (e.g., *do, don't*), then all modalities are included in the sponsor except for the outer epistemic: John didn't *want to try to ski* but Mary did [e].

When *non-modal* matrix verbs occur in the sponsor clause, they are usually excluded from the sponsor, e.g., *refused to* in (5). However, if that non-modal matrix verb is scoped over by modality, then it is included in the sponsor, e.g., *ask* in (6).

(5) Mylonas refused to *reveal* the ransom paid for his release, although police sources did [e].

(6) He wanted to *ask* me to come but didn't [e].

This rather coarse-grained rule set has been under refinement since the evaluation reported here, with results to be reported separately.

Although resolving reference to a text string does not address the issue of sloppy coreference – one of the most studied aspects of ellipsis in theoretical linguistics – it has a significant practical benefit: the system requires as input only a raw (not manually corrected) syntactic parse along with several readily sharable word lists and rule sets. In other words, this system is ready to be incorporated into application systems, or could easily be reimplemented in other environments.

The Engines

The system takes as input our indexed version of the Gigaword corpus (Graff and Cieri 2003) and consecutively launches the following processing engines.

Modal scope ellipsis detector. This engine extracts elliptical contexts of interest by finding verbal modal elements (e.g., *can, won't*) that are directly followed by a hard discourse break, defined as a period, semi-colon or colon. This maximally lightweight detection strategy achieves high precision at low cost, but at the expense of recall.

Syntactic analyzer. The sentences that are hypothesized to contain modal-scope ellipsis are processed by the Stanford dependency parser (version Stanford-corenlp-2012-01-08; de Marneffe, MacCartney and Manning 2006), whose output serves as input to our rule-based resolution strategy.

Elliptical "phrasal" detector. The first ellipsis resolution strategy seeks configurations that can be treated like open patterns, i.e., multi-word expressions (phrasals) containing a combination of fixed and variable elements, such as *wherever NP *can*, *all the N (that) NP *can,* and *everything (that) NP *can* (the notation **can* refers to any inflectional form of the verb, such as *could, cannot,* etc.). All of the phrasals treated so far function as adverbials, so a simple resolution strategy – selecting the preceding verbal element as the sponsor – works quite well. For example, sentence (7) matched the pattern *what NP *can* and correctly indicated that the sponsor was *say*.

(7) Vincent Schmid, the vicar of the cathedral, said prayer and music would *say* what words could not [e].

Simple parallel configuration detectors. A key linguistic insight guiding our approach is that *syntactically simple* contexts that show specific types of *structural parallelism* provide high predictive power without the need to resort to deep semantic or pragmatic analysis. Let us define what we mean by *structural parallelism* and *syntactic simplicity*.

Structural Parallelism. Each structurally parallel context contains an ellipsis clause directly preceded by a conjunct that is syntactically connected to in one of three ways that can be loosely described as exhibiting syntactic parallelism: coordination (be it of clauses or verb phrases), parataxis, and certain types of main/adverbial clause pairs. The Stanford parser calls these dependencies "conj" (8), "parataxis" (9) and "advcl" (10), respectively. For each example, the key dependencies are indicated, and, for readability, the word-number indices output by the parser are included only when needed for disambiguation. As mentioned earlier, only the leftmost element of the sponsor is italicized.

(8) They *kept* attacking and we didn't [e].
conj(kept, did)

(9) In one walk, they *used* poles; in the other, they didn't [e].
parataxis(used,did)

(10) Oakland's Barry Zito took the victory, *hurling* eight innings as Santana did [e].
advcl(hurling, did)

The motivation for exploiting syntactic parallelism to predict ellipsis resolution derives both from the well-documented linguistic effects of parallelism (e.g., Goodall, 2009; Hobbs and Kehler 1998; McShane 2000, 2005; McShane and Beale, Submitted) and successful results of our experiments.

Syntactic simplicity. The predictive power of parallel configurations decreases significantly if the conjuncts – particularly the first – contain relative or subordinate clauses because such elements provide additional candidate sponsors for the elided verb phrase. For example, if we rewrite corpus-attested (11) such that the first clause includes several embedded clauses, as in (11a), it becomes necessary to carry out sophisticated reasoning about the world to determine which preceding verb is the leftmost constituent of the sponsor: *autographed? broken? playing? left?*

(11) Bush *autographed* a steady stream of memorabilia, but Garciaparra, his hand in a cast, could not [e].

(11a) Bush *autographed* a steady stream of memorabilia, but Garciaparra, who had broken his wrist playing soccer the day before he left for the conference, could not [e].

We operationalized the notion of "simple parallel" configurations in terms of the output of the Stanford parser. Configurations are deemed "simple parallel" if they contain: (a) exactly one instance of a *conj, advcl* or *parataxis* dependency – which we call "whitelisted" dependencies – that links the modal element licensing the ellipsis with an element from the sponsor clause; (b) no instances of a "blacklisted dependency" – i.e., *ccomp, rcmod, dep* or *complm* – which indicate various types of embedded verbal structures that complicate matters by offering competing candidate sponsors; (c) one or more instances of a "graylisted" dependency, defined as an *xcomp* or *aux* dependency that takes as its arguments matrix and/or main verbs from the sponsor clause. The parse for sentence (12) includes one whitelisted dependency, ***conj**(wanted, did),* and three graylisted dependencies: ***xcomp**(wanted, try), **xcomp**(try, start), **xcomp**(start, juggle)*.

(12) John wanted to *try* to start to juggle and did [e].

Although graylisted dependencies make sponsor clauses more complicated, and although contexts containing them can offer more than one reasonable resolution for the ellipsis (in (12), either *try* or *start* could be interpreted as the left-hand constituent of the sponsor), the added complications can be managed by the rule sets for treating sponsor-clause matrix verbs described above.

We configured 6 search patterns to detect and resolve simple parallel configurations. Three of them do not permit modal verbs in the sponsor clause, though non-modal matrix verbs are permitted: Simple Conjunction NoMod (8, above), Simple Parataxis NoMod (9), Simple Adverbial Clause NoMod (10). The other three simple parallel configurations contain one or more modal verbs in the sponsor clause, as well as, optionally, non-modal matrix verbs: Simple Conjunction Mod (13), Simple Adverbial Clause Mod (14) and Simple Parataxis Mod (no corpus examples of the latter were found).

(13) Fans in other countries could *apply* on the Internet, but fans in the host nation could not [e].
Whitelisted: **conj**(apply-6, could-17).
Graylisted: **aux**(apply-6, could-5)

(14) I tried to *drive* hard as I could [e].
Whitelisted: **advcl**(drive, could)
Graylisted: **xcomp**(tried, drive)

Although "simple parallel" configurations do not represent a large percentage of elliptical examples in the corpus, when they are found, they offer quite confident predictions about the ellipsis sponsor.

Less Simple Parallel Configuration Detector. In an attempt to increase recall, we experimented with relaxing the notion of "simplicity" in syntactically parallel configurations by permitting the parse to contain a single extra whitelisted dependency (*conj, advcl, parataxis*) or a single blacklisted dependency (*ccomp, rcomd, dep, complm*). The goal was to operationalize the intuition that, although having many extra candidate sponsors due to syntactic complexity would dramatically reduce the predictive power of the engine, having just one additional point of complexity might not be too detrimental, especially if we could classify the most benign types of additional dependencies. For example, in (15), the whitelisted dependency that establishes the parallelism used to resolve the ellipsis is ***advcl****(exploded, did)*, and the "extra" whitelisted dependency that can be ignored is ***parataxis****(thrown, exploded)*.

(15) Two of the devices were thrown over a security fence; one *exploded*, while the other did not [e].

A combination of introspection and corpus analysis suggested that the most benign additional dependency was *parataxis.* So for this precision-oriented system run, we included just two "less simple parallel" configurations: Simple Adverbial Clause Mod and Simple Adverbial Clause NoMod, each with one additional parataxis dependency.

Non-Parallel, Closest Modality Detector. This engine does not seek structural parallelism; instead, it walks back through the text and seeks the first modal verb, hypothesizing that it might be "paired" with the ellipsis licensor and therefore point to the sponsor. (16) is an example correctly resolved by this strategy: *could* is the most recent modal verb, which is stripped from the sponsor due to the modality-correlation rules described above.

(16) Rivard and Cosworth's Steve Miller both said they could *build* an Indy only engine similar to Mercedes for next year if they must [e].

This concludes the description of the current state of our ellipsis detection and resolution system. The output of the system is a set of examples that the system (a) believes include modal-scope ellipsis prior to a hard discourse break, and (b) believes it can, with high confidence, resolve, with "resolution" defined as indicating the leftmost verbal element of the sponsor. We next describe an evaluation exercise that both validates the utility of the approach and provides pattern-level (rather than corpus-level) confidence metrics for systems that might employ the results.

Evaluation, Analysis, Lessons Learned

Our original research goals included learning as much as possible about realizations of modal-scope ellipsis in the Gigaword corpus; building a robust system to treat those phenomena; and detailing a plan of work to expand the system and cover residual phenomena. To assess our progress, we carried out a combined evaluation and glassbox analysis of the linguistically motivated hypotheses described above and the engines we used to implement them. Evaluation was carried out by a graduate student judge whose decisions were selectively checked by a senior developer. Evaluation results are presented below. We are confident that launching our engines on other comparable corpora would yield similar results.

Evaluating elliptical "phrasals". Of the 18,298 Gigaword sentences that were identified as having modal-scope ellipsis before a hard discourse break, the system detected 486 unique examples that it believed matched one of our phrasal configurations. We randomly selected about 1/3 of those – 161 examples – to be manually evaluated. Among those, there were 8 detection errors (nominals, like *a can*, were mistaken for verbs), yielding a detection accuracy of

95%. Of the 153 actually elliptical examples, 147 were resolved correctly, yielding a resolution accuracy of 96%.

Evaluating parallel and non-parallel configurations. The system detected 1,814 unique Gigaword inputs that it believed matched one of the "simple parallel", "less simple parallel" or "non-parallel" configurations. All of these were manually evaluated. There were 24 ellipsis detection errors, yielding a detection accuracy of 98%. The actually elliptical 1,790 examples were evaluated for resolution accuracy, with results as shown in Table 1.

Table 1. Evaluation results for non-phrasal configurations. Fractions were rounded down. P1/P2: phases 1 and 2 of testing.

Pattern	Exs.	MatrixVerb Err. (P1/P2)	Misc. Err.	Precision (P1/P2)
Simp.Conj.Mod	238	2 / 2	1	98 / 98 %
Simp.Conj.No.Mod	364	83 / 48	15	73 / 83 %
Simp.Adv.Cl.Mod	132	3 / 3	2	96 / 96 %
Simp.Adv.Cl.No.Mod	662	16 / 10	53	89 / 90 %
Simp.Para.No.Mod	54	0	4	92 / 92 %
Less.Simp.Parallel	15	0	2	80 / 80 %
NonParallel	325	8 / 5	60	79 / 80 %

Column 2 indicates the number of examples that matched each configuration listed in Column 1. Columns 3 and 4 show error tallies, with 4 indicating miscellaneous errors and 3 reflecting errors in which a matrix verb should have been stripped from the sponsor but was not. In phase 1 of testing (P1), we had not yet incorporated the non-modal-matrix-verb stripping rules described above, so, e.g., the system selected *refused* rather than *reveal* as the head of the sponsor for example (5). The phase 2 system (P2), by contrast, incorporated that rule and correctly treated this, and many other, cases of non-modal matrix verbs. We cite this example of stepwise system enhancement to show the natural progression of glassbox testing/evaluation in a knowledge-based system. Naturally, the phase 2 evaluation statistics reflect only the introduction of globally applicable rules, not instance-level tweaks that would have undercut the overall validity of the evaluation results.

One source of residual errors related to processing sponsor-clause matrix verbs results from the parser's misinterpretation of certain matrix + complement constructions. Details aside, for some matrix verbs, if the logical subject of the complement clause is a full NP (*She urged* ***the crowd*** *to disperse peacefully*), the parser fails to establish the expected xcomp dependency between the matrix verb (*urged*) and its complement (*disperse*). However, if that full NP is replaced by a pronoun (*She urged* ***them*** *to disperse peacefully*), the expected xcomp dependency is returned. We have experimented with working around this error by automatically replacing full NPs in such configurations by pronouns, but the workaround is not yet reliable enough for inclusion in the system. As a sidebar, we consider the incorporation of such workarounds a necessary response to real-world constraints, and a much better option than building systems that rely on prerequisites that cannot be automatically supplied.

Space does not permit exhaustive reporting of our error tracking, but let us mention just three types of errors that are easily remedied. First, our modality correlation rules, which initially only applied to verbs, should be extended to adjectives and nouns that convey modal meanings, such as *opportunities* (which conveys potential modality) in (17). Second, topicalization strategies in the sponsor sentence, such as *NP copula* ***to*** ***V_{INFIN}*** (18) can be exploited in new detection/resolution patterns.

(17) I had my opportunities to *break* and I didn't [e].

(18) The first target was to *qualify* for the Asian Cup and we did [e].

Third, we hypothesize that the performance of the non-parallel pattern will improve dramatically if we introduce two types of linguistically motivated constraints: (1) the pattern will accept only semantically strong correlations between the ellipsis licensing modal and the most recently encountered modal: e.g., *tried...failed* and *could…couldn't* will be treated, but *believe...want to* will not; (2) the system will evaluate only the most recent verbal structure to determine if it contains such a modal element, it will not be permitted to skip over candidate sponsors in search of a modal. Both of these constraints will reduce recall in favor of precision. Further research is needed to determine how to judiciously expand recall without unacceptable losses in precision.

In contrast to these more easily remedied errors is the hard residue of difficult phenomena, such as contexts in which the actual sponsor is located outside of the window of coreference. Typical evaluation setups would manually remove such instances, but we include them in our results to accurately portray how the system can perform when launched, unaided, on open text.

System enhancement is continuing in two stages. We are currently: implementing the three easy fixes just described; enhancing the ellipsis-detection process to significantly increase recall; expanding the inventory of phrasal configurations and the lists of non-modal matrix verbs; improving our matrix stripping rules to more precisely cover specific combinations of modal and non-modal matrix verbs; investigating how our current rule sets apply to VP ellipsis sponsored by a non-modal verb (*I went because she advised me to [e]*); and redefining "resolution" to include the entire span of the sponsor rather than just the leftmost element. Our practical strategies will continue to rely on minimal preprocessing requirements, making the resulting approach and system widely applicable.

The next stage of system enhancement – moving toward the type of full interpretation discussed in McShane et al. 2012b – will rely on the semantic and pragmatic analysis engines of the OntoSem environment. For example, we will seek other types of non-parallel configurations that offer strong predictive power, including those that leverage semantic constraints; we will seek linguistically motivated methods of loosening the definition of "simplicity" in simple parallel configurations to increase system recall; and we will investigate how reference resolution for subjects can inform the more precise reconstruction of elided categories in terms of instance vs. type event coreference and strict vs. sloppy argument coreference. We must emphasize that all of these 2nd stage enhancements will involve analysis and algorithms that pertain not only to elided categories but to overt referring expressions as well. For example, the overt referring expressions *do it* and *do that*, as well as some uses of pronominal *this, that* and *it,* refer to propositions, and detecting their sponsors involves the same reasoning as detecting the sponsors for elided VPs. Based on our past work in resolving overt and elided *nominal* referring expressions (McShane and Nirenburg 2013; McShane and Beale, Submitted), we hypothesize that the resolution algorithms for overt and elided v*erbal* referring expressions will largely overlap.

Discussion

The system described here treats an important subset of elliptical phenomena – elided scopes of modality – in fully automatic mode. The approach is primarily knowledge-based in that it leverages linguistic observations involving syntax, semantics and discourse, but it uses only readily available resources: the Stanford dependency parser, several word lists, and lightweight engines that match input sentences to ellipsis resolution patterns and then resolve the ellipsis.

Our main claims are: (1) that treating *some* instances of difficult linguistic phenomena is preferable to avoiding difficult phenomena wholesale; (2) that configuring systems that can work in fully automatic mode – from detecting which phenomena will be treated, to treating them and approximating the system's confidence in each answer – is preferable to configuring systems that require manually prepared input, such as annotated corpora; and (3) that incorporating knowledge-based methods into intelligent agent systems in practical, readily extensible ways offers the best balance between near-term utility and long-term progress.

The main motivation for our overall program of research is to develop cognitively plausible intelligent agents. Work reported in this paper is no exception, which is made manifest along several lines. First, the ellipsis detection and resolution strategies incorporate well-attested linguistic phenomena, such as syntactic parallelism effects, a well-motivated inventory of modal meanings and their lexical realizations, and predictable semantic correlations between modal meanings (*try…fail; can…cannot*).

Second, the approach reflects a practical interpretation of the principle of least effort, which arguably underpins much of human behavior. It is unlikely that a person engages in deep semantic and pragmatic analysis to resolve the ellipsis in a sentence like *Mary napped but I didn't [e]*; instead, this is more likely involves a cognitively cheaper, ready-made resolution strategy that can be implemented, for example, through analogical reasoning. We implement this strategy by first leveraging cheap, high-confidence "phrasals" (lexicalized patterns), which are ready-made, combined solutions to ellipsis detection and resolution problems. Next, we invoke more compositional but still very high-confidence patterns; and so on, covering ever more cases but with ever higher processing demands, lower precision and lower confidence. This progression from "ready-made solution" to "fully compositional analysis" is prominent across all text understanding modules of OntoSem.

Third, the necessity of resolving elided categories is absolutely central to artificial intelligent agent systems, particularly if they are to participate in dialog applications, since dialog prominently features a broad array of elliptical phenomena. In other words, ellipsis cannot wait until the next pendulum swing of preferred methods in NLP brings explanatory modeling back into the focus of the field.

Fourth, in order for intelligent agents to collaborate effectively with people, they need to be able to independently evaluate their confidence in their own language understanding. The approach described here offers detailed, pattern-level confidence metrics, which should be of much greater utility to intelligent agents than a single, corpus-level measure.

Finally, this work will contribute directly to the functioning of agents in the OntoAgent architecture, as they attempt to arrive at the full meaning of every utterance as input to subsequent reasoning and action.

Acknowledgments

This research was supported in part by Grant N00014-09-1-1029 from the U.S. Office of Naval Research. All opinions and findings expressed in this material are those of the authors and do not necessarily reflect the views of the Office of Naval Research. Thanks to Ben Johnson for his work on an early implementation of this system.

References

Byron, D. 2004. Resolving Pronominal Reference to Abstract Entities. PhD Dissertation, University of Rochester. Available as Technical Report 815, University of Rochester Computer Science Department.

de Marneffe, M., MacCartney, B., and Manning, C. D. 2006. Generating Typed Dependency Parses from Phrase Structure Parses. In Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC).

Fiengo, R., and May, R. 1994. *Indices and Identity*. Cambridge, Mass.: The MIT Press.

Graff, D., and Cieri, C. 2003. English Gigaword. Linguistic Data Consortium. Philadelphia.

Goodall, G. 2009. *Parallel Structures in Syntax: Coordination, Causatives and Restructuring*. Cambridge University Press.

Hardt, D. 1997. An Empirical Approach to VP Ellipsis. *Computational Linguistics* 23(4): 525-541.

Hobbs, J., and Kehler, A. 1998. A Theory of Parallelism and the Case of VP Ellipsis. In Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics (ACL-98), pp. 394-401.

Lee, H., Chang, A., Peirsman, Y., Chambers, N., Surdeanu, M., and Jurafsky, D. 2013. Deterministic Coreference Resolution Based on Entity-Centric, Precision-Ranked Rules. *Computational Linguistics* 39(4): 885-916.

McShane, M. 2005. *A Theory of Ellipsis*. Oxford University Press.

McShane, M. 2000. Hierarchies of Parallelism in Elliptical Polish Structures. *Journal of Slavic Linguistics* 8: 83-117.

McShane, M., and Beale, S. Submitted. Reference Resolution for Intelligent Agents: A Focus on Confidence. Submitted to *Journal of Language Modelling*.

McShane, M., Beale, S., Nirenburg, S., Jarrell, B., and Fantry G. 2012a. Inconsistency as a Diagnostic Tool in a Society of Intelligent Agents. *Artificial Intelligence in Medicine*, 55(3): 137-48.

McShane, M., and Nirenburg, S. 2012. A Knowledge Representation Language for Natural Language Processing, Simulation and Reasoning. *International Journal of Semantic Computing* 6(1).

McShane, M., and Nirenburg, S. 2013. Use of Ontology, Lexicon and Fact Repository for Reference Resolution in Ontological Semantics. In Oltramari, A., Vossen, P., Qin, L., and Hovy, E. (eds.), *New Trends of Research in Ontologies and Lexical Resources*. Springer, pp 157-185.

McShane, M., Nirenburg, S., and Beale, S. Submitted. The OntoSem Text Understanding System. Submitted to *Advances in Cognitive Architectures.*

McShane, M., Nirenburg, S., Beale, S., Johnson, B. 2012b. Resolving Elided Scopes of Modality in OntoAgent. *Advances in Cognitive Systems,* vol. 2, Dec. 2012.

Nirenburg, S., McShane, M., and Beale, S. 2008. A Simulated Physiological/Cognitive "Double Agent". In Beal, J., Bello, P., Cassimatis, N., Coen, M., and Winston, P. (eds.), Papers from the AAAI Fall Symposium, "Naturally Inspired Cognitive Architectures," Washington, D.C., Nov. 7-9. AAAI technical report FS-08-06, Menlo Park, CA: AAAI Press.

Nirenburg, S., and Raskin, V. 2004. *Ontological Semantics*. Cambridge, Mass.: The MIT Press.

Pang, B., and Lee, L. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval* 2: 1–135.

Spenader, J., and Hendriks, P. 2006. Combining Multiple Information Sources for Ellipsis. *Research on Language and Computation* 4(4): 327-333.