

The Interplay of Language Processing, Reasoning and Decision-Making in Cognitive Computing

Sergei Nirenburg and Marjorie McShane

Rensselaer Polytechnic Institute, Troy, NY, 12180, USA
zavedomo@gmail.com, margemc34@gmail.com

Abstract. Integrating language processing, reasoning and decision making is a prerequisite to any breakthroughs in cognitive computing. This paper discusses historical attitudes that have worked against such integration, then describes a cognitive architecture called OntoAgent that illustrates both the feasibility and the payoffs of pursuing integration. Examples are drawn from the Maryland Virtual Patient prototype application, which offers medical trainees the opportunity to diagnose and treat a cohort of cognitively modeled virtual patients that are capable of language processing, reasoning, learning, decision making and simulated action.

Keywords: natural language processing, intelligent agents, reasoning, cognitive architecture

1 The Analytic Approach and Its Consequences

Since at least the times of Descartes, the scientific method has become more or less synonymous with the analytical approach, whereby a phenomenon or process is decomposed into contributing facets or components. The general idea is that, after each such component has been sufficiently studied independently, there would follow a synthesis step that would result in a comprehensive explanation of the phenomenon or process. A well-known example of the application of the analytical approach is the tenet of the autonomy of syntax in theoretical linguistics, which has been widely adopted by – and has strongly influenced – the field of computational linguistics.

The analytical approach makes good sense because it is well nigh impossible to expect to account for all the facets of a complex phenomenon simultaneously and at a consistent grain size of description. But it comes with a cost: it artificially constrains the purview of theories and the scope of models, and it has unwittingly fostered indefinite postponement of the all-important synthesis step.

Within the field of cognitive modeling, the analytical approach is evident in the traditionally defined pipeline of agent functionalities: *perception*, *reasoning*, *action*. Under this view, natural language understanding is subsumed under *perception* and natural language generation is subsumed under *action*. And,

although it is well understood that human-level natural language processing will require extensive reasoning, decision-making and learning, these *language-oriented* manifestations have not been on the agenda of the respective research communities.

This paper explores our efforts at synthesizing, within the OntoAgent cognitive architecture, the treatment of natural language understanding with the treatment of reasoning, decision-making, learning, and non-linguistic perception. After introducing the historical landscape and briefly describing our cognitive architecture, we will illustrate this synthesis using the following four simple premises as motivation: (1) While natural language understanding provides input to reasoning, reasoning, in turn, supports the process of natural language understanding. (2) While natural language understanding provides input to decision-making, decision-making, in turn, is an integral part of the process of natural language understanding. (3) Not only does NLP support an agent’s learning about the language and world, learning about language and the world is needed for the agent to engage in human-level language understanding. (4) Agent memory can be populated not only by results of natural language understanding; it can be populated by output from other channels of perception.¹ Although we have been working toward actualizing these aspects of integration for some time, this paper represents our first attempt to generalize about how this effort contributes to the goal of *synthesis at an early stage* in the development of human-inspired intelligent agents.

1.1 Natural Language Processing Within Cognitive Modeling, Historically

Although natural language processing (NLP) requires reasoning and reasoning requires NLP, the mainstream NLP and reasoning communities diverged early in the history of AI, with the gap continuing to widen over time. Mainstream NLP focuses on the shallow analysis of uninterpreted text strings (see [7] for a balanced overview of the state of the art), avoiding language problems that require anything beyond statistically-oriented computation. Automatic reasoning systems, for their part, typically rely on hand-crafted, unambiguous input, assuming that language problems will eventually be solved by external processors (see, e.g., [10]). This division between NLP and reasoning was recognized already in the 1950s by Bar Hillel:

“...The evaluation of arguments presented in a natural language should have been one of the major worries of logic since its beginnings. However, [...] the actual development of formal logic took a different course. It seems that [...] the almost general attitude of all formal logicians was to regard such an evaluation process as a two-stage affair. In the first stage, the original language formulation had to be rephrased, without

¹ Results of reasoning can also populate the agent memory. We do not develop this topic in this paper.

loss, in a normalized idiom, while in the second stage, these normalized formulations would be put through the grindstone of the formal logic evaluator. [...] Without substantial progress in the first stage even the incredible progress made by mathematical logic in our time will not help us much in solving our total problem.” ([2], pp. 202-203)

If we substitute the more modern terms “knowledge representation language” for “normalized idiom,” and “reasoning system” for “formal logic evaluator,” it becomes clear that the state of affairs described by Bar Hillel has only modestly changed over the last half-century. That is, in spite of the important work on semantically-oriented NLP by groups led by Schank [24], Wilks [29], Woods [30], Allen [1], Schubert [27] and others, the “evaluation of arguments presented in a natural language” has yet to garner intensive, widespread attention. Instead, the center of gravity for mainstream NLP has been on statistical tools applied to text strings in large text corpora.

As mentioned earlier, in the study of cognitive architectures, it is customary to modularize agent functionalities. The most coarse-grained categorization was presented above: *perception, reasoning, action*. Various more fine-grained classifications have also been put forth, such as the one found in Langley et al.’s 2009 survey article [9]. They describe nine capabilities that any good cognitive system must have: 1) recognition and categorization; 2) decision making and choice; 3) perception and situation assessment; 4) prediction and monitoring; 5) problem solving and planning; 6) reasoning and belief maintenance; 7) execution and action; 8) interaction and communication; and 9) remembering, reflection and learning. The authors primarily subsume NLP under “interaction and communication” but acknowledge that it involves other aspects of cognition as well. They recognize the lack of fundamental integration of NLP with other aspects of agent cognition, stating, “Although natural language processing has been demonstrated within some architectures, few intelligent systems have combined this with the ability to communicate about their own decisions, plans, and other cognitive activities in a general manner.” Indeed, of the 18 representative architectures briefly described in the Appendix, only two – SOAR [12] and GLAIR [26] – are overtly credited with involving NLP; and one, ACT-R, is credited indirectly by reference to applied work on tutoring [8] within its framework. Of all of the cross-modular influences, one that has been particularly well explored is the interaction between NLP and planning: e.g., the pioneering work of Cohen, Levesque and Perrault (see [5] and [22]) demonstrated the utility of approaching NLP tasks in terms of AI-style planning, and planning is a first-order concern in the field natural language generation [23].

The reason why NLP has been addressed rather peripherally and not in depth in the field of cognitive modeling is because it is a very hard problem; more precisely, it is a collection of many hard problems. However, satisfactory solutions to these problems can only be expected if they are tackled as part of overall agent functioning. In other words, at least outside the realm of “low-hanging fruit” applications, NLP must be integrated with other aspects of simulated

agent cognition. This is the hypothesis pursued in the OntoAgent program of research and development, to which we now turn.

2 OntoAgent and the Maryland Virtual Patient Application

OntoAgent can be described as both a cognitive architecture and a knowledge environment: it has the same goals as traditional cognitive architectures, but, unlike many other architectures, it also stresses descriptive and system-building work aimed at creating a non-toy knowledge substrate to support the functioning of agents in applications [15]. OntoAgent has three core static knowledge resources: the ontology, the fact repository and the lexicon. The **ontology** is a knowledge base of descriptions that contains knowledge about types of objects, events, relations that link them and attributes that describe them. It also contains the scripts that support agent simulation. The current version of the OntoAgent ontology contains about 9,000 concepts described by an average of 16 properties each. Some agents are not endowed with the full ontology, largely to simulate individual differences in knowledge of specialized domains. The **fact repository** contains remembered instances of ontological concepts and their property values. Naturally, each agent has its own fact repository to reflect its individual simulated experience. The **lexicon** describes approximately 30,000 word senses of English, both syntactically and semantically, with semantic descriptions written in the ontological metalanguage.

Language understanding in OntoAgent means automatically translating ambiguous and often elliptical natural language inputs into an unambiguous, ontologically-grounded metalanguage suited for reasoning. This translation is carried out using heuristic evidence that relies primarily on information in the lexicon and ontology (for a description of the analysis process, see [20]). The analyzer produces text meaning representations like the one shown in Table 1, which represents the meaning of the English sentence *Dolores has severe chest pain*. Concepts are shown in small caps, and numerical suffixes differentiate instances. The similarity of concept names to English words is solely to support manual knowledge acquisition: the meaning of a concept is defined as its inventory of property values.²

Ontologically-grounded text meaning representations facilitate expectation-driven reasoning during both language processing and decision-making. This is because the ontology contains more information about each *type* of object and event than is known about each textually attested *instance*. For example, whereas the concept instance PAIN-23 in our example includes fillers for the properties INTENSITY and LOCATION, the ontology includes many more properties of PAIN. One is PAIN-CAUSE-TYPE, whose literal fillers are *nociceptive*, *neuropathic* and *psychogenic*. When the concept PAIN is activated (i.e., used in a text meaning

² The same ontology can be used for representing the meaning of utterances in any language, given an ontological semantic lexicon for that language.

PAIN-23	EXPERIENCER	HUMAN-37
	INTENSITY	.8
	LOCATION	CHEST-BODY-PART-14
	textstring	“has”
	from-sense	have-v18 (a phrasal entry for <i>have...pain</i>)
HUMAN-37	EXPERIENCER-OF	PAIN-23
	HAS-NAME	“Dolores”
	HAS-GENDER	female
	textstring	“Dolores”
	from-sense	*personal-name*
CHEST-BODY-PART-14	LOCATION-OF	PAIN-23
	textstring	“chest”
	from-sense	chest-n1

Table 1. Text meaning representation for *Dolores has severe chest pain*.

representation), the agent may expect further dialog to include references to as-yet “unused” properties. So, given the subsequent input, *The pain was due to tissue damage.*, the polysemous phrases *due to* and *tissue damage* will be analyzed as the highly specific property-value pair PAIN-CAUSE-TYPE NOCICEPTIVE using expectation-driven reasoning.

The application that has served as the substrate for validating our approach to agent modeling is Maryland Virtual Patient (MVP). MVP is a prototype clinician training application that features a cohort of cognitively modeled virtual patients that can be diagnosed and treated by human trainees in open-ended simulations [14] [19]. Virtual patients are comprised of linked physiological and cognitive simulations. Physiologically, virtual patients change over time and in response to interventions by the user. Cognitively, virtual patients can engage in dialog with the user, make decisions about their health care and lifestyle, learn and remember new information, and carry out simulated action. The bridge between physiology and cognition is *interoception*, defined as the perception of one’s bodily signals. We model interoception as one of two channels of perception, the other being natural language understanding. Both interoception and natural language understanding generate identical meaning representations that are used to populate the agent’s fact repository (cf. Section 3.4).

Along with virtual patients, the MVP environment features a tutoring agent, which can provide context-sensitive guidance to the trainee. As described in [16] and [17], the same knowledge substrate used by the tutor could be used to support the functioning of an advisor to practicing clinicians.

3 Four Examples of Integrating NLP With Other Agent Functionalities

This section briefly describes the previously introduced four points of integration of NLP with other aspects of agent cognition in OntoAgent.

3.1 NLU Supports Reasoning and Reasoning Supports NLU

The fact that natural language understanding (NLU) supports reasoning is well-attested in the cognitive systems literature: after all, NLU is typically viewed as a type of perception, which precedes reasoning and action. For example, in the MVP application, once a virtual patient has understood the meaning of (i.e., generated a text meaning representation for) “What brings you here?” it must:

1. Detect the goal(s) that the interlocutor is pursuing in uttering the dialog turn.
2. Integrate the results of its analysis of text meaning and speaker goal into its memory.
3. Decide to generate an instance of a “Be-a-Cooperative-Conversationalist” goal and add it to its active goal agenda.
4. Prioritize goal instances on the agenda (in this case, the above instance will be prioritized).
5. Select a plan to pursue to attain this goal (in this case, the plan will be to carry out a verbal action).
6. Decide on content of the verbal action to be produced (in this case, this will involve checking its memory for recent symptoms).
7. Generate an English sentence that realizes the above content by outputting, for example, “I’ve been having difficulty swallowing.”

In short, natural language understanding launches a cascade of other agent reasoning functions.

But just as NLU supports reasoning, so must reasoning be brought to bear for NLU, since the challenges presented by natural language are formidable: lexical ambiguity, referential ambiguity, idiomaticity, ellipsis, indirect speech acts, non-literal language, unexpected input and more. These challenges have been underplayed in the past 20 years, as mainstream NLP has chosen to focus on those linguistic phenomena that are most amenable to supervised machine learning. For example, there have been significant efforts toward detecting (but not semantically interpreting) multi-word expressions [25]; resolving the simpler cases of textual coreference [11]; and selecting which relations hold between uninterpreted nominals in nominal compounds [28]. However, as long as NLP is approached as the manipulation of uninterpreted textual strings, its results will not be sufficient to support human-level reasoning by intelligent agents.

Within OntoAgent, by contrast, we do not shun the difficult problems posed by natural language. Addressing such problems naturally requires knowledge-based, reasoning-intensive methods [20] [13] [21]. Of course, the goal of fully understanding all language phenomena in open text will not be achieved overnight – all of the contributing algorithms and the knowledge bases they rely on require

long-term, iterative improvement. However, even setting the goal of language *understanding* shifts the perspective away from isolated NLP modules and toward the integration of natural language into overall agent cognition.

3.2 NLU Supports Decision-Making and Decision-Making is Needed for NLU

As we just saw, NLU supports reasoning and reasoning is a prerequisite to decision-making, so the statement “NLU supports decision-making” should be self-evident. However, the reverse dependency – i.e., decision-making in support of NLU – has received little attention. This is in large part because NLU has been recently treated as a monolithic process that results in a singularly right or wrong answer. However, this orientation fails to account for the fact that normal, unedited, natural language use does not consist of exclusively “clean” utterances that can be understood – even by people – with 100% precision. Instead, language is littered with false starts, infelicitous ellipses, intentional and unintentional vagueness, unnecessary detail, incomprehensibly formulated thoughts, and so on. For this reason, NLU is better modeled as a multi-stage process after each stage of which the agent asks itself *Have I understood enough to proceed to reasoning (and action)?* Decisions about “enough” will spare agents from endlessly pursuing ever deeper language analysis.

Let us consider a few examples in which “enough” is achieved at different stages of processing.

- *The basic text meaning representation is sufficient.* The basic text meaning representation is sufficient to support reasoning when an agent is faced with a direct question (*Do you have chest pain?*) or a direct command (*Please tell me your symptoms.*). In these cases, the basic text meaning representation includes an instance of a REQUEST-INFO (request information) or REQUEST-ACTION event, which is sufficient for the agent to generate a verbal action in response.
- *Indirect speech act detection is needed.* When the basic text meaning representation does not include a direct request for information or action, the agent attempts to determine whether an indirect speech act was used, as is the case in the following: *I’d like to know if you ever have chest pain* (indirect question), *I think that surgery is your best option* (indirect request for action). In such instances, the result of the agent’s reasoning is recorded in a so-called *extended* text meaning representation, which includes the initially masked REQUEST-INFO or REQUEST-ACTION concept.
- *Not all input needs to be fully understood.* As mentioned above, it is not uncommon for natural language input to include words, phrases and even whole sentences that are functionally superfluous. Such is the case, for example, when the speaker precedes a request for information or action by a long preface: *I know we’ve talked about a lot of things related to your past and current symptom profile, but what I’d really like to know at this point in time is, do you have chest pain?* Even if the agent cannot not confidently

disambiguate every lexeme preceding the question, it can still answer the question and hold up its end of the dialog interaction. We are not suggesting that it is optimal for an agent to fail to fully understanding something; however, we are suggesting that if the goal is to build useful intelligent agents in the near- and mid-term, teaching them to focus on actionable aspects of utterances is well-motivated.

- *Clarification from the human collaborator is needed.* In some cases, the agent might fail to understand a necessary portion of an input and therefore be blocked from subsequent reasoning. This can happen, for example, if an agent is asked a question about an unknown word/concept: *Do you feel pain in the area of your lower esophageal sphincter?* In such cases, the agent must make the decision to pursue learning during language processing itself, as described further in Section 3.3.
- *The agent decides reason about the speaker’s goal.* In many cases, people respond not only to the direct meaning of a question or request, but to their understanding of the speaker’s goal in uttering it. For example, the following dialog turns are quite natural: (a) *“Where are your keys?” “You can’t borrow my car.”* (b) *“I have a stomach ache.” “You’re going to school.”* (c) *“Can you run and fetch me a screwdriver?” “This knife will work just as well.”* In each of these cases, responding to the direct meaning would have also been appropriate, but the interlocutor considers these responses more efficient. Whether a speaker will respond to the direct meaning of an utterance or to his/her interpretation of the speaker’s goal is a function of the interlocutor’s personality traits, understanding of the situation, the relationship between the speaker and interlocutor, and so on.

Consider an example in which five different virtual patients (VPs), who present to the doctor with the symptom of coughing, respond to the question *“Have you been traveling lately?”* in different ways for different reasons.³

- VP1: No, I haven’t been anywhere that might have made me sick.
- VP2: Yes, I was on a crowded plane last week.
- VP3: No.
- VP4: Yes, I drove to Washington to visit my sister.
- VP5: No. Why are you asking?

VP1 and VP2 have an inventory of personality traits and physical and mental states that compel them to hypothesize about the goal the speaker is pursuing. Since they are reporting to the doctor with a complaint, and since they have not yet been diagnosed, they hypothesize that the doctor’s goal is diagnosis. They try to figure out – using ontological search – how COUGH is connected with TRAVEL-EVENT. The ontologies of these two VPs are the same, and they include the information that COUGH can be CAUSED-BY INFLUENZA, that INFLUENZA is a COMMUNICABLE-DISEASE, and that COMMUNICABLE-DISEASE can be CAUSED-BY AIRPLANE-TRAVEL, BUS-TRAVEL, TRAIN-TRAVEL (the latter is a simplification of a much longer causal chain that involves being located in crowded spaces). So, they understand that their cough might be caused by something they encountered

³ See [18] for further discussion of agent parameterization.

during these types of travel. The two patients differ, however, with respect to their fact repositories (i.e., memories of past experiences). VP1 does not have any recorded memories of relevant travel events – i.e., travel in an airplane, bus or train – so it responds ‘no’. The remainder of the utterance serves as a trace that it is responding to the goal behind the question, not to its literal meaning. VP2, by contrast, has a different fact repository: it recently traveled on a plane. So it answers positively, and the elaboration of its response serves as a trace that it responded to the speaker’s goal as well. VP3, for its part, decides to bypass goal-oriented mindreading and to respond to the direct meaning of the question.⁴ The decision not to read extra meaning into the question is quite natural for this agent since it has no recent travel events in its fact repository – i.e., nothing to trigger decision-making about whether or not an event is relevant.

As for VP4, it is unclear from its utterance whether or not it attempted to reason about the speaker’s goal. Either it did attempt to and failed, or it simply answered the surface meaning of the question. If it failed, then the likely reason is that its ontology lacks the necessary link between coughing and certain types of travel.

Finally, VP5 decides not to mindread but attempts to learn what goal the user is pursuing in asking this question, reflecting a knowledge-seeking character trait.

In sum, it would be counterproductive to assume that all agents must at all times understand every input to the same depth and with the same precision. Some inputs defy precise analysis, even by humans; some inputs provide more detail than is needed for the sufficient functioning of an agent in an application; and many inputs can be analyzed and responded to differently by agents with different knowledge, histories, traits and states. All of this variability is best supported by opening up the many stages of the language understanding process and giving the agent an option to make decisions about how to proceed following each of those stages.

3.3 NLP Supports Learning and Learning Supports NLP

Learning in OntoAgent is understood as the agent’s ability to populate, modify and use its knowledge bases – ontology, lexicon and fact/belief repository – over the period of its functioning (its “lifespan”). As an illustration of agent learning, let us look at the MVP dialog excerpt shown in Table 2. At the time point we observe, the human agent is interviewing the virtual patient and recommends that she have an EGD. This particular virtual patient does not know the word EGD or the concept EGD – they are absent from her lexicon and ontology. So, she must engage in further dialog to learn the needed information (NLP supports

⁴ We can only know for sure that this agent did not undertake goal-related mindreading by looking at the trace of system processing. VP1 could also have generated just ‘No’ as its response.

learning), and she must learn the information in order to carry out further dialog on this topic (learning supports NLP). Traces of system functioning (which can be seen in dynamic “under the hood” panels of the interface) are presented in italics. The fact that the patient asks questions about the procedure is due to its character traits: it requires extensive information before making a decision.

MD: **I suggest having an EGD, which is a diagnostic procedure.**
 VP: *I record EGD as a child of DIAGNOSTIC-PROCEDURE in my ontology.*
I record EGD as a noun in my lexicon, mapped to EGD.
I don't know enough about its properties to agree.
I need to ask questions about risk and pain.
How risky is it?
 MD: **It's not risky at all.**
 VP: *Add RISK 0 to ontological specification of EGD.*
Is it painful?
 MD: **It's only a little uncomfortable.**
 VP: *Add PAIN .1 to ontological specification of EGD.*
I can tolerate that amount of risk and pain.
OK, I'll agree to that.

Table 2. Dialog between the doctor (MD) and the virtual patient (VP). Dialog turns are in boldface, whereas traces of VP decision-making are in italics.

3.4 Memory Population from Natural Language Understanding or Other Perception Modes

Perception is the primary source of agent learning. As mentioned above, OntoAgents have two channels of perception: language understanding and simulated interoception. Perception via both of these channels results in identical meaning representations, formulated in the ontological metalanguage. Earlier (Table 1) we saw how text analysis results in such a meaning representation; let us now consider the analogous process of interoception.

During physiological simulation, when certain property values reach a given threshold, they trigger the instantiation of a symptom. For example, when a virtual patient's esophagus is sufficiently inflamed, this triggers the symptom of heartburn. Table 3 shows how the ontological representation of a mild esophageal inflammation, which is generated by the physiological simulation of gastroesophageal reflux disease, automatically generates the symptom *mild, occasional heartburn*. What is noteworthy for this discussion is that the ontological representation of this perceived symptom is the same as the meaning representation for the sentence, “I am experiencing mild, occasional heartburn.” So, whether the agent experiences this symptom or is told it has this symptom (for whatever reason the latter might happen), it will be stored to memory in the same form and can support the same subsequent reasoning.

INFLAMMATION-1	LOCATION	ESOPHAGUS-1
	EFFECT	CHEST-PAIN-1
ESOPHAGUS-1	LOCATION-OF	INFLAMMATION-1
	PART-OF-OBJECT	PATIENT-1
CHEST-PAIN-1	INTENSITY	.1
	FREQUENCY	.1
	EXPERIENCER	PATIENT-1
	CAUSED-BY	INFLAMMATION-1
PATIENT-1	EXPERIENCER-OF	CHEST-PAIN-1
	HAS-OBJECT-AS-PART	ESOPHAGUS-1

Table 3. Ontological representation for esophageal inflammation generating the symptom of mild, occasional chest pain. Intensity and frequency of chest pain are measured in the abstract scale 0,1.

4 Final Thoughts

We have selectively illustrated the tight integration of language processing with reasoning, decision making, learning and alternative channels of perception in the OntoAgent environment. Addressing this integration is, we believe, a prerequisite to any breakthroughs in semantic and cognitive computing. In turn, success in this undertaking depends on the availability of a rather long list of capabilities and resources that are currently only partially available. In our opinion, the most crucial of these is sufficiently fine-grained knowledge resources, such ontologies and lexicons. Therefore, R&D in knowledge acquisition must assume an ever growing importance.

The current trend in computational linguistics is “supply-side” – pursuing broad coverage of acquired resources (at the expense of the depth of description of knowledge elements) using machine learning techniques supervised as lightly as possible. By contrast, the long-continuing trend in the fields of cognitive architectures and application systems that deploy language capabilities is still predominantly “demand-side” – devoting just as much effort to resource development as is minimally needed, with “minimally needed” being defined differently in the different paradigms. In cognitive architectures, language capabilities must be minimally sufficient to test the reasoning algorithms and representational formalisms, whereas in application systems, language capabilities must be minimally sufficient to achieve user acceptance.

In the framework of OntoAgent – and work that preceded it in the paradigm of Ontological Semantics – we have been developing a hybrid approach that seeks a better balance between supply-side and demand-side approaches, as well as between the breadth and depth of acquired knowledge. It is our hope that recent and current knowledge acquisition efforts – e.g., [3], [4], [6] – bring in

new insights and results that will permit the research community to succeed in building adequate knowledge resources for cognitive computing.

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.

References

1. Allen, J., Ferguson, G., Stent, A. An architecture for more realistic conversational systems. Proceedings of the Conference on Intelligent User Interfaces, 1-8 (2001).
2. Bar Hillel, Y. Aspects of Language. Jerusalem: Magnes (1970).
3. Bobrow, D., Condoravdi, C., Karttunen, L, Zaenen, A. Learning by reading: normalizing complex linguistic structures onto a knowledge representation. Proceedings of the AAAI Spring Symposium on Learning by Reading and Learning to Read (2009).
4. Clark, P., Harrison, P. Large-scale extraction and use of knowledge from text. Proceedings of the Fifth International Conference on Knowledge Capture, ACM, 153-160 (2009).
5. Cohen, P. R., Levesque, H. J. Rational interaction as the basis for communication. In Cohen, P. R., Morgan, J., Pollack, M. E. (Eds.), Intentions in Communication. Cambridge, MA: MIT Press (1990).
6. Hahn, U., Marko, K. G. Ontology and lexicon evolution by text understanding. Proceedings of the ECAI 2002 Workshop on Machine Learning and Natural Language Processing for Ontology Engineering (2002).
7. Jurafsky, D., Martin, J. H. Speech and Language Processing: An Introduction to Natural Language Processing, Speech Recognition, and Computational Linguistics. 2nd edition. Prentice-Hall (2009).
8. Koedinger, K. R., Anderson, J. R., Hadley, W. H., Mark, M. A. Intelligent tutoring goes to school in the big city. International Journal of Artificial Intelligence in Education, 8, 30-43 (1997).
9. Langley, P., Laird, J. E., Rogers, S. Cognitive architectures: Research issues and challenges. Cognitive Systems Research, 10, 141-160 (2009).
10. Langley, P., Meadows, B., Gabaldon, A., Heald, R. Abductive understanding of dialogues about joint activities. Interaction Studies 15(3), 426-454 (2014).
11. Lee, H., Chang, A., Peirsman, Y., Chambers, N., Surdeanu, M., Jurafsky, D. Deterministic coreference resolution based on entity-centric, precision-ranked rules. Computational Linguistics 39(4), 885-916 (2013).
12. Lewis, R. An architecturally-based theory of human sentence comprehension. PhD Thesis. Carnegie Mellon University. CMU-CS-93-226 (1993).
13. McShane, M. Reference resolution challenges for an intelligent agent: The need for knowledge. IEEE Intelligent Systems, 24 (4), 47-58 (2009).
14. McShane, M., Jarrell, B., Fantry, G., Nirenburg, S., Beale, S., Johnson, B. Revealing the conceptual substrate of biomedical cognitive models to the wider community. In Westwood, J. D., Haluck, R. S., et al. (Eds.) Medicine Meets Virtual Reality 16, 281 - 286. Amsterdam, Netherlands: IOS Press (2008).

15. McShane, M., Nirenburg, S. A knowledge representation language for natural language processing, simulation and reasoning. *International Journal of Semantic Computing* 6 (2012).
16. McShane, M., Beale, S., Nirenburg, S., Jarrell, B., Fantry, G. Inconsistency as diagnostic tool in a society of intelligent agents. *Artificial Intelligence in Medicine (AIIM)*, 55(3), 137-48 (2012).
17. McShane, M., Nirenburg, S., Jarrell, B. Modeling decision-making biases. *Biologically-Inspired Cognitive Architectures (BICA) Journal*, 3, 39-50 (2013).
18. McShane, M. Parameterizing mental model ascription across intelligent agents *Interaction Studies* 15(3), 404-425 (2014).
19. Nirenburg, S., McShane, M., Beale, S. A simulated physiological/cognitive “double agent”. In Beal, J., Bello, P., Cassimatis, N., Coen, M., 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 (2008).
20. McShane, M., Nirenburg, S., Beale, S. *Language Understanding With Ontological Semantics*. *Advances in Cognitive Systems* (forthcoming).
21. McShane, M., Beale, S., Babkin, P. Nominal compound interpretation by intelligent agents. *Linguistic Issues in Language Technology (LiLT)*, 10 (2014).
22. Perrault, C. R., and Allen, J. F. A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics*, 6, 167-182 (1980).
23. Reiter, E. Natural language generation. In: Clark, A., Fox, C., Lappin, S. (Eds.) *The Handbook of Computational Linguistics and Natural Language Processing*. Wiley-Blackwell, 574-598 (2010)
24. Schank, R., Riesbeck, C. *Inside computer understanding*. Hillsdale, NJ: Erlbaum (1981).
25. Schone, P., Jurafsky, D. Is knowledge-free induction of multiword unit dictionary headwords a solved problem? *Proceedings of Empirical Methods in Natural Language Processing*, Pittsburgh, PA. (2001).
26. Shapiro, S. C. , Ismail, H. O. Anchoring in a grounded layered architecture with integrated reasoning. *Robotics and Autonomous Systems*, 43(2-3), 97-108 (2003).
27. Schubert, L., Hwang, C. H. Episodic logic meets Little Red Riding Hood: A comprehensive, natural representation for language understanding. In Iwanska, L., Shapiro, S. (Eds.), *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*. Menlo Park, CA and Cambridge, MA: MIT/AAAI Press (2000).
28. Tratz, S., Hovy, E. A taxonomy, dataset, and classifier for automatic noun compound interpretation. *Association for Computational Linguistics* (2010).
29. Wilks, Y., Fass, D. Preference semantics: A family history. *Computing and Mathematics with Applications* 23(2) (1992).
30. Woods, W. A. *Procedural semantics as a theory of meaning*. Research Report No. 4627. Cambridge, MA: BBN (1981).