

Modeling abduction within Ontological Semantics

Max Petrenko

Purdue University

CERIAS

mpetrenk@purdue.edu

Victor Raskin

Purdue University

CERIAS

vraskin@purdue.edu

Abstract

The paper presents an attempt to emulate abductive forms of reasoning within the framework of Ontological Semantics. Abduction is understood as a meaning-based and goal-oriented inference-construction mechanism, which functions as a disambiguation and interpretation tool. First, the concept of abduction is discussed based on current literature. Particularly, the meaning-based and goal-oriented nature of abduction is considered. The Ontological Semantics NLP system is then introduced. It is argued in the paper that the rich knowledge resources and versatile computational mechanisms of Ontological Semantics-supported parser allow it to emulate, in computational implementations, the processing natural language as it is employed by humans. An example of elliptic input with non-verbalized case-roles is presented, and the major stages of processing are outlined in a step-by-step fashion, with screenshots of the knowledge resources and parser. A rule is introduced which captures some dependencies among the semantic properties of the elliptic and non-elliptic segments of input.

Introduction

The present paper shows how ontological Semantics, a multi-modular knowledge-based NLP system, can emulate abduction-based disambiguation and interpretation process, intrinsic to human semantic competence, for use in computer systems.

For an NLP system to function at the level of human-like complexity entails being capable of processing input containing non-verbalized segments. Elliptic information is extracted effortlessly by humans due to an extensive (and potentially infinite) world model and a set of dynamic semantic mechanisms.

The world model, which includes background knowledge, immediate knowledge and situation-dependent goals, supplies an abducting human with a pool of potential solutions. Semantic mechanisms assist selecting among most salient solutions in order to reconstruct elliptic information. This meaning-based and goal-oriented mechanism of inference-construction has been defined as abduction (Hobbs et al., 1994; Walton, 2004; Gabbay and Woods, 2005; Aliseda, 2006).

Ontological Semantics (Nirenburg and Raskin, 2004) has all the necessary potential

for emulating abduction. Most importantly, Ontological Semantics-based applications operate on a richly specified and hierarchically modeled knowledge resource base (i.e. world model), which is indispensable for any semantics-motivated NLP enterprise (on the notion and relevance of “doing semantics semantically”, as opposed to *du jour* methods of NLP skirting the semantics prerequisite, see Hempelmann and Raskin, 2008).

The following sections will discuss the notion of abduction, introduce Ontological Semantics system and will show how human-based abduction is reflected in the Ontosem-based processing module powered by inference rules.

1. Ontological Semantics and the abductive nature of human competence

The sections will provide a theoretical discussion of abduction in natural language processing and Ontosem-based parser and identify how the former is paralleled by the latter in its structure.

1.1. Abduction in natural language processing

Stemming from the works by Peirce (Peirce, 1955), abduction is defined in current theories as a two-stage reasoning process of generating a set of hypotheses and selecting the most suitable yet potentially defeasible one. The precise nature of the two stages of abductive reasoning, the understanding of the hypothesis “suitability” criteria varies across theories thus branching the current field into explanatory, i.e. algorithm- and rule-based (Peng and Reggia, 1990; Hobbs et al., 1994; Flach, 2000; Wang, 2000) and

probabilistic (Pearl, 2000) approaches to the hypothesis evaluation (for an extensive review of the two approaches see Gabbay and Woods, 2005, and Thaggard, 2000 for an attempt to synthesize the two).

Most generally, the dynamics of abductive reasoning could be schematized as follows:

Stage I: hypotheses generation

- 1) A requires explanation;
- 2) No immediate explanation available;
- 3) N explanations are generated;

Stage II: hypothesis selection

- 4) Explanation B meets criteria X
- 5) Explanation B is selected;

(a rough generalization derived from Walton, 2004; Aliseda, 2006; Gabbay and Woods, 2005, see also Peirce, 1955; Niiniluoto, 1998; Allan, 2001; Attardo, 2003; Paavola, 2004 for a syllogism-type structure of abduction).

The meaning-based and goal-oriented nature of abduction is described in Gabbay and Woods (2005): cognitive agents (humans and larger entities) abduce by deriving a most plausible and effective explanation which fulfils immediate goals pertaining to the specific situation (see the section discussing the “instrumental” nature of abductive explanation: p.81). Along the same lines, Aliseda (2006) argues that an abductive solution is always embedded in a specific knowledge framework (i.e. “background theory”).

The following section will introduce the Ontological Semantics system. It will then be argued that Ontological Semantics reflects the two-stage structure of abduction

as well as its meaning-based and goal-oriented nature.

1.2. Ontological Semantics per se

The Ontological Semantics system features a powerful set of static knowledge resources and dynamic algorithms for sentence parsing, semantic processing and sense disambiguation. The current version of Ontological Semantics comprises:

- An ontology of ca. 10,000 language-independent concepts with hierarchical structure and properties-inheritance mechanism. The root concept ALL branches into three major classes covering objects, events and properties (Figure 1). The ontology in Ontological Semantics plays similar role as the world model in human

competence: each of the three branches splits into several sub-branches covering objects, events and properties (including relations, literal, scalar attributes, modality values, etc.) with respective properties which are propagated from the top-level concept to every child concept through the inheritance mechanism. For a detailed review of the branches with more useful illustrations see (Hempelman and Raskin, 2008).

The richly specified system of interrelations among the ontological concepts sets Ontological Semantics apart from other versions of (pseudo-) ontologies based on controlled vocabularies and described in Orbst (2007) and discussed in Raskin et al. (2008); Hempelman and Raskin (2008).

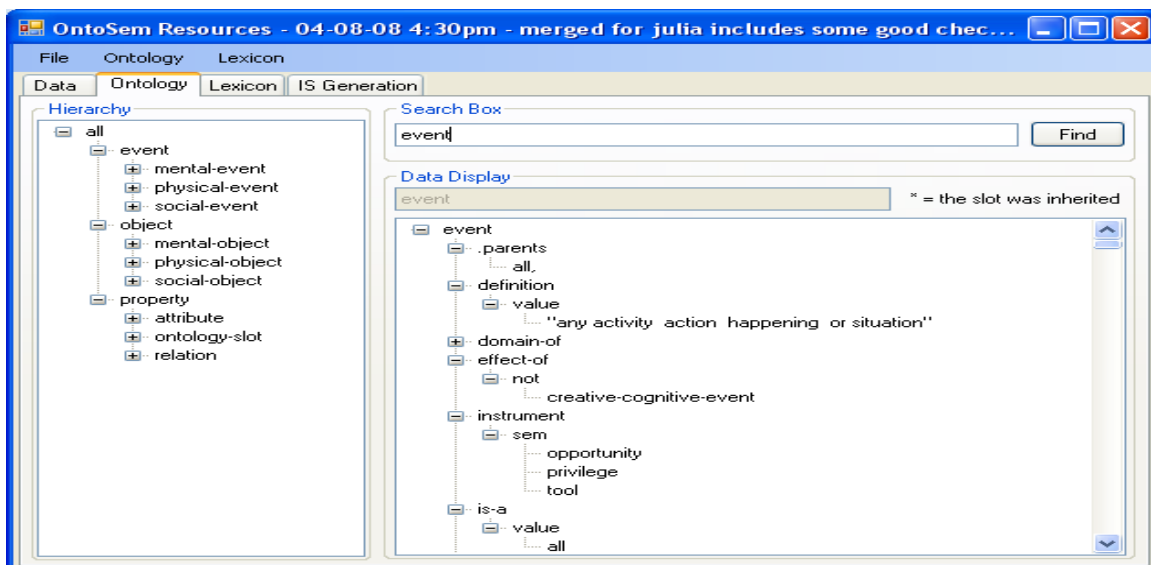


Figure 1: the three top-level concept branches with the EVENT branch properties

- Several language-specific lexicons including English with ca. 120,000 senses, and several smaller ones; Onomastica – several language-specific dictionaries of proper names with ca. 25,000 senses. Each entry of a lexicon contains syntactic structure (grammatical and combinatorial features) and semantic structure (a concept instantiated by the entry with properties indicated, where necessary) (see Figure 2).

- A text-meaning-representation (TMR) format, which captures the meaning of single- and multi-clausal sentences input.
 - A fact repository, which stores instances of TMR's for immediate goals of input processing;
 - A dynamic Ontological Semantics-based parsing module which analyzes the input by identifying the clauses, instances of concepts, events and their case-roles.

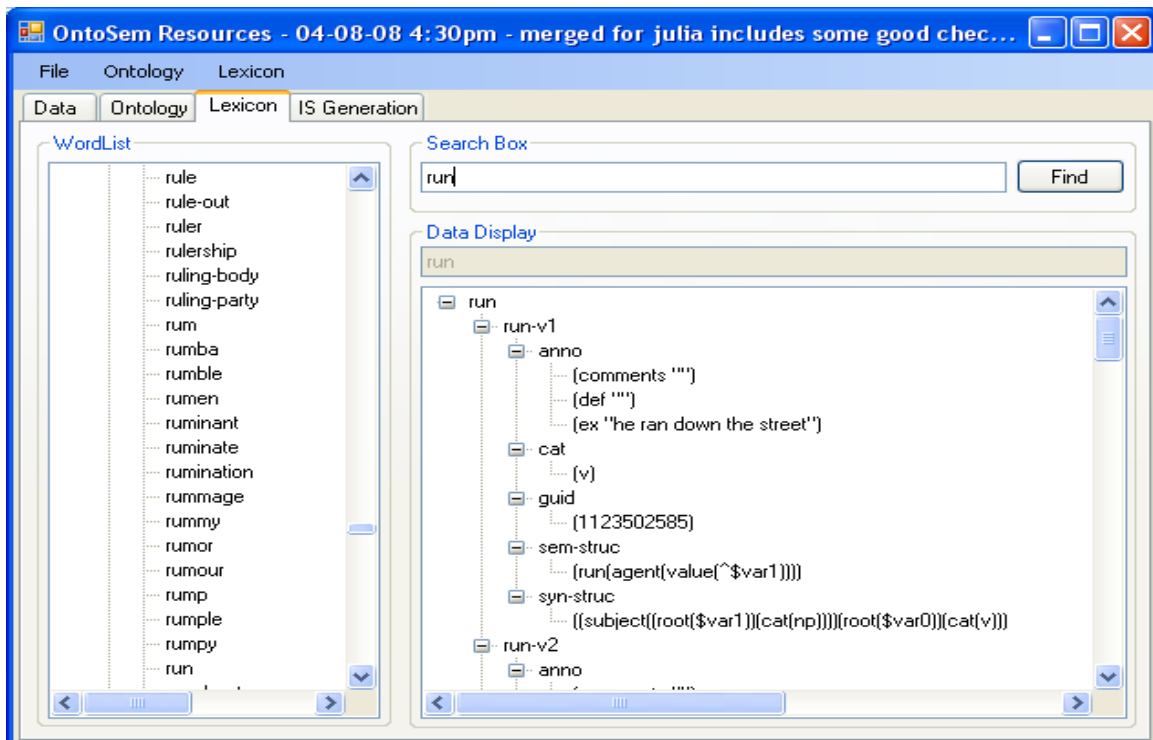


Figure 2: A lexical entry for the intransitive sense “run-v1” in the lexicon

After the pre-semantic processing (part-of-speech and morphological tagging and clause-breaking) of the input is complete,

the Ontological Semantics-based parser takes every meaningful entry in a clause and looks up for the underlying concept starting

from the head units (events) and matching the remaining ones (objects and properties) with the case roles of the identified event. Non-verbalized fillers for major case roles (agent, theme, time, location) can be identified based on the default properties of the ontological concepts (e.g. HUMAN as a default agent of ACTIVE-COGNITIVE-EVENT). When the input is elliptic and no default fillers are available, the inference module makes an attempt to derive ones.

All the instances of concepts and their constraining properties are recorded in the fact repository. The fact repository plays a crucial role in Ontosem-based applications detecting contradictions. An example of the Ontosem-based parser flagging conflicting case-role fillers of the same event across clauses is analyzed in Raskin et al (2008).

2.3 Abduction in Ontological Semantics

The rich knowledge resources and versatile computational mechanisms of Ontological Semantics allow the Ontosem-based parser to emulate abduction in its human-like complexity.

The meaning-based and goal-oriented nature of abduction, as well as its two-staged structure as described in section 2.1 can be observed in the way the parser operates.

First, the system proceeds in both top-bottom and bottom-up directions. Input processing starts with an inquiry to the knowledge base (Ontology, Onomasticon), where the lexical entry in question is assigned a concept. After major concepts (primarily events) are identified and the initial TMR is submitted to the fact repository, all additional input is processed

based on the acquired knowledge with co-reference resolved and event case-roles filled by default.

The system thus follows the same pattern as abducting humans do: project the background knowledge on the context of the given situation so that your hypothesis search would narrow down substantially. This allows humans to select effectively and quickly most salient solutions from a potentially unlimited repository.

Second, the two stages of abductive reasoning are replicated in Ontological Semantics-based processing as the multiplicity of potential case-role fillers which is dynamically constrained to the most salient one(s).

Informed by the ontology, the parsing module is typically presented with several fillers for a required case-role. The state of having a broad pool of potential case-role fillers thus parallels that of an abducting human deriving a set of plausible explanations of a fact.

The Ontological Semantics-based parser then employs a set of inferential rules trying to single out suitable case-role filler. This state of evaluating the suitability of candidate filler and selecting one is similar to that of an abducting human quickly determining the inference most salient to the given situation.

Similar to human competence, the parsing module is robust and difficult to baffle: non-verbalized case-role fillers are reconstructed based on by-default properties of ontological concepts and context-dependent case-role dependencies. Input containing unattested

entities is processed through assigning a respective case-role to the entity in question and sliding the ontology down to its most immediate ontological parent. Nirenburg and Raskin (2004) describe an example of processing a sentence *Fred locked the door with the key* (see section 8.4.3). The LOCK-EVENT identified in the entry *lock* contains an instrument case role, which *key* fills in based, on the one hand, on the pre-processing of the preposition *with* and, on the other hand, on the ontology-provided ARTIFACT filler of the instrument case-role of LOCK-EVENT. The entity *key*, although unrecorded in the lexicon, would still be processed meaningfully.

The section below will demonstrate how elliptic input (the benchmark of maturity of an NLP application), can be processed by the Ontosem-based parser powered by inference rules. It will be shown that non-verbalized non-default fillers for events can be reconstructed across-clauses.

The reader will be guided through the steps of the Ontosem-based processing of an example presenting three semantically interrelated events two of which contain non-verbalized case-role fillers.

3 Ontological Semantics-emulated

abduction: ellipsis processing

The selected example,

May 5, 2005: A small device made of training grenades stuffed with black powder is thrown at a Manhattan building about

3:55 a.m. Small-scale damage, but no injuries,

contains three clauses two of which instantiate two events with implicit fillers for the properties of time, location, instrument and theme.

At the first stage, the clause-parsing module will break the sentences into three clauses based on three events THROW, DESTROY, INJURE instantiated in the lexemes “throw”, “damage”, “injure”. The lexeme “made”, which maps onto the concept CREATE-ARTIFACT in the ontology, by virtue of having a Past Participle form without a be-copula and following a noun, will be syntactically identified as modifying the noun “device” and thus will not require a separate clause.

By the same token, “stuffed”, which maps onto the concept CHANGE-EVENT (with the effect specified as OBJECT containing material), will not form a clause of its own but will be identified as a modifier of the noun “grenades”. The results of the clause-breaking stage are displayed on Figure 3.

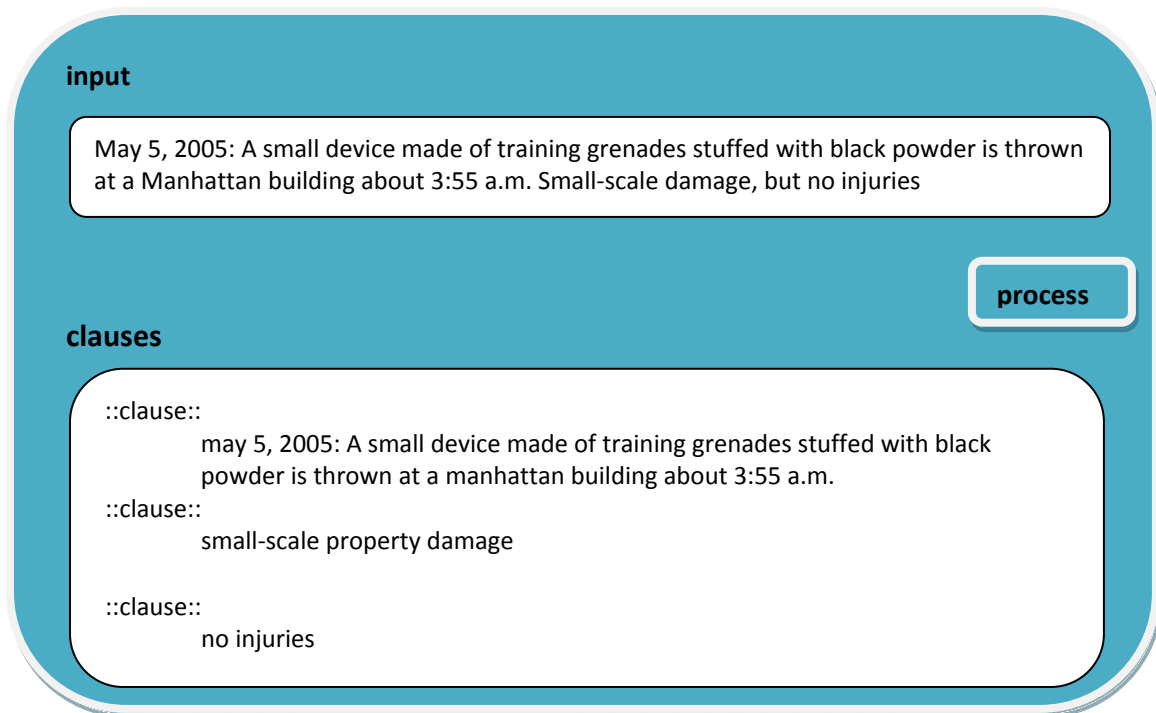


Figure 3: clause-breaking module

After the clause-breaking phase, the input is checked for the head of the proposition. The process of finding a head is event-driven: it operates on the principle “most proposition heads are events” and proceeds from checking the case roles for a given event and matching its constraints with the properties of other entities given in the input. In case no explicit events are present, objects are prioritized over properties (see Nirenburg and Raskin, 2004, section 8.2.1 for a detailed description of algorithms and principles for establishing propositional structure). Three events, instantiated in “thrown”, “damage”, and “injuries” are

given in the input: THROW, DESTROY, and INJURE, each forming a separate clause.

The system will then attempt to establish semantic relations among the events across clauses by trying every following concept as filler for the properties of every preceding one. Thus, the effect property of THROW (constrained by HAND-GRENADE as its theme) would have INJURE and DESTROY filling its effect case-role slots.

After relations among the clause-forming events have been established, the head event THROW is checked for its case roles, and possible fillers are identified based on the

input. Those properties, for which there is no verbalized filler, are assigned default fillers. The location property of THROW will be filled by the concept DISTRICT constrained by the has-name property with “Manhattan” as its value. The theme case role will be filled by PHYSICAL-OBJECT constrained by made-of(sem(hand-granade)) and size attribute. The time property will have several fillers listed, as provided in the input: the timestamp, the month and the year. The destination case role will be filled by BUILDING. Thus the following TMR will be derived by the parser (see Figure 4). DESTROY has unspecified (not provided by ontology) and non-verbalized fillers for its time, location and theme properties. INJURE has unspecified fillers for time and location, properties, and the theme case-role by-default filled by ANIMATE. The disambiguation is achieved by employing an inference rule:

IF
 (1) (E1(effect(E2)));
 (2) E1 is MOVEMENT-EVENT;
 (3) No explicit properties of E2 are available,

THEN
 (E1(time)) = (E2(time));
 (E1(location)) = (E2(location));
 (E1(destination)) = (E2(theme)),

Simply put, the rule stipulates that that in elliptic input (i.e. when no explicit or reconstructable by default fillers are available) two events (one of which is a MOVEMENT-EVENT) linked through the “effect” property would share an agent,

time, location and have equal fillers for the destination and theme properties.

An algorithm has been developed based on the rule:

1. Identify concepts. Events present?
 Yes – Proceed to 2; No – terminate
2. E1 is MOVEMENT-EVENT?
 Yes – proceed to 3; No – terminate.
3. Identify relations among events.
 E1(effect(E2))?
 Yes – proceed to 4; No – terminate;
4. Check E1 properties: identify fillers for time; location; destination;
5. Check E2 properties: identify fillers for time; location; theme;
 Fillers missing – proceed to 6;
 Fillers present – terminate.
6. Use:
 filler for (E1(time)) as filler for (E2(time));
 filler for (E1(location)) as filler for (E2(location));
 filler for (E1(destination)) as filler for (E2(theme));

As a result of the processing, the fillers for time and location properties of THROW will be transferred to their counterparts for DESTROY and INJURE. The destination property filler for THROW will fill the theme case-role slot for DESTROY. Figure 5 illustrates the extended TMR generated by the parser after the inference rule has been applied (marked in red, blue and green are identical fillers for the three events).

Thus, Ontological Semantics-based application has all the potential to process elliptic input in its human-like complexity.



Figure 4: TMR derived based through standard processing

May 5, 2005: A small device made of training grenades stuffed with black powder is thrown at a Manhattan building about 3:55 a.m. Small-scale damage, but no injuries

process

Events filled

```
(throw
  (time(value(month(has-name(value("may"))))
        (year(cardinality(value(2005))))))
  (time-hour(value(3)))
  (location(sem(district(has-name(value("manhattan"))))))
  (theme(sem(physical-object(made-of(sem(grenade(size(value(<0.5))))))))
  (destination(sem(building)))
  (effect(sem
    (destroy
      (time(value(month(has-name(value("may"))))
            (year(cardinality(value(2005))))))
      (time-hour(value(3)))
      (location(sem(district(has-name(value("manhattan"))))))
      (theme(sem(building)))
      (instrument(sem(physical-object
                    (made-of(sem(grenade(size(value(<0.5))))))))
      (epistemic(value(<0.5)))
    (injure
      (time(value(month(has-name(value("may"))))
            (year(cardinality(value(2005))))))
      (time-hour(value(3)))
      (location(sem(district(has-name(value("manhattan"))))))
      (theme(sem(animate)))
      (instrument(sem(physical-object
                    (made-of(sem(grenade(size(value(<0.5))))))))
      (epistemic(value(0))))))
```

Figure 5: extended TMR as a result of the inference rule application.

Acknowledgment

This research was supported in part by the National Science Foundation Grant #0428554 to Purdue University. The authors are grateful to the anonymous reviewers for their comments.

References

- Allan, K. 2001. *Natural Language Semantics*. Malden, MA: Blackwell Publishers.
- Aliseda, A. 2006. *Abductive reasoning: logical investigations into discovery and explanation*. National Autonomous University of Mexico. 2006.
- Attardo, S. 2003. On the Nature of Rationality in (Neo-Gricean) Pragmatics. *International Journal of Pragmatics*. 14. 3-20. (Special issue on Neo-Gricean pragmatics, edited by Ken Turner).
- Flach, P. 2000. On the Logic of Hypothesis Generation. In: Flach, Peter A., and Kakas, Antonis C. (eds.). 2000 *Abduction and induction : essays on their relation and integration*. Kluwer Academic Publishers, Dordrecht, Boston, MA. 117-129.
- Hempelmann, C.F. and V. Raskin. 2008. Semantic Search: Content Vs. formalism. In: *Proceedings of Langtech 2008*, Rome, Italy. http://www.langtech.it/en/technical_program/technical_program.htm (full paper).
- Hobbs, J.R., Stickel, M.E., Appelt, D.E., and Paul Martin. 1994. Interpretation as Abduction. In: Pereira, F.C.N. and Barbara J. Grosz (eds.) *Natural Language Processing*. MIT Press. 69-142.
- Gabbay, D. M., and J. Woods. 2005. The Reach of Abduction. *Insight and Trial. A Practical Logic of Cognitive Systems*, 2. Elsevier Science, Amsterdam, the Netherlands.
- Niiniluoto, I. 1998. Defending Abduction. *Philosophy of Science*, Vol. 66, Supplement. Proceedings of the 1998 Biennial Meetings of the Philosophy of Science Association. Part I: Contributed Papers. (Sep., 1999): S436-S451.
- Nirenburg, S., and V. Raskin 2004. *Ontological Semantics*. Cambridge, MA: MIT Press. A prepublication draft, chapter by chapter, can be found on ontological semantics.com.
- Obrst, L. 2007. Ontology and Ontologies: Why it and they Matter to the Intelligence Community. Proceedings of the Second International Ontology for the Intelligence Community Conference. OIC-2007. Columbia, MD. November 28-29.
- Paavola, Sami. 2001. Abduction as a Logic and Methodology of Discovery: The importance of strategies. To appear in *Foundations of Science*. <http://www.helsinki.fi/science/commens/papers/abductionstrategies.html>
- Peng, Y., and Reggia, J. A. 1990. *Abductive Inference Models for Diagnostic Problem-*

- Solving*. Springer-Verlag, New York and Berlin.
- Pearl, J. 2000. *Causality: models, reasoning, and inference*. Cambridge University Press. Cambridge, UK; New York.
- Peirce, Ch. S. 1955. *Philosophical Writings of Peirce*. New York: Dover. 1960-1966. *Collected Papers*. Charles Hartshorne and Paul Weiss (eds.): Harvard University Press. 8 Vols., Cambridge, MA.
- Raskin, V., S. Nirenburg, I. Nirenburg, C. F. Hempelmann, and K. E. Triezenberg 2003. The genesis of a script for bankruptcy in ontological semantics. In: G. Hirst and S. Nirenburg (eds.), *Proceedings of the Text Meaning Workshop, HLT/NAACL 2003: Human Language Technology and North American Chapter of the Association of Computational Linguistics Conference*. ACL: Edmonton, Alberta, Canada, May 31.
- Raskin, V., C. F. Hempelmann, and K. E. Triezenberg 2004. Semantic Forensics: An Application of Ontological Semantics to Information Assurance. In: G. Hirst and S. Nirenburg (eds.), *Proceedings of the Second Workshop on Text Meaning and Interpretation, ACL, Barcelona, Spain, July 25-26, 2004*. ACL Publication. Edmonton, Alberta, Canada.
- Raskin, V., C. F. Hempelmann, and K. E. Triezenberg 2008. Ontological Semantic Forensics: Meaning-based Deception Detection. *Meaning computation*, 1(forthcoming).
- Thaggard, Paul. 2000. *Coherence in Thought and Action*. The MIT Press.
- Walton, D. 2004. *Abductive Reasoning*. University of Alabama Press.
- Wang, P. 2000. Unified Inference in Extended Syllogism. In: Flach, Peter A., and Kakas, Antonis C. (eds.). 2000. *Abduction and induction : essays on their relation and integration*. Kluwer Academic Publishers, Dordrecht, Boston, MA. 117-129.