Delivering Comprehensive Knowledge of the World to the Computer:

How to make the computer understand meaning

Victor Raskin
Linguistics & CERIAS
Purdue University
West Lafayette, Indiana, USA
Email: vraskin@purdue.edu

Julia Taylor Rayz
Computer and Information Technology
Purdue University
West Lafayette, Indiana, USA
Email: jtaylor1@purdue.edu

Abstract—**Ontological Semantic Technology is a mature theory that we have developed for 30 and 14 years, respectively to the co-authors, and it has been demonstrated to represent the meaning of text well. Designed first as an engineering ontology for natural language computer applications, it became a paradigm for theoretical, as well as computational semantics. This paper claims that it is also a legitimate approach to a much larger cognitive task of representing the human knowledge of the world in the computer, without which text understanding is unattainable.**

Keywords-ontology; ontological semantics; semantic theory; representing computer knowledge of the world

I. Introduction: Representing Meaning

The main purpose of this position paper is to demonstrate how a comprehensive ontology, extended from forming the basis and substance of natural language semantics, can and must be extended to the basis and substance of human knowledge that is computerized in applications. The former task was established in Ontological Semantics developed in the late 1980s through 1990s.

In a parallel and unrelated development, the same decades saw the growing presence of non-linguistic natural language processing that became dominant. It makes sense now, for a number of reasons, that computational semantics, the linguistic field, and natural language processing, the non-linguistic "techie" field, would combine their efforts in the near future. This paper, coming out of the former field addresses researchers in both areas: it shows the computational semanticists how the approach reaches into the cognitive area of modeling human knowledge of reality for the computer; and it informs the 'techies' of the existing and developing effort outside of machine learning, neural networks, deep learning, word embedments, and such.

Section II traces the genesis and trajectory of ontological semantic, thus establishing its basic tenets. Section III deals briefly with the other approach, commenting only on what matters for this paper. Section IV explains how the ontological semantics paradigm progressed from the basis of limited computational semantic application to that of semantic theory in general, Section 5 introduces the ontological plane as that substance in which semantic symbols are interpreted in. Immediately, Section VI shows how the approach can, then, handle the modeling of comprehensive human knowledge for the computer. Section VII is the brief conclusion.

II. Mission of Ontological Semantics

Ontological Semantics originated in the late 1980s and was developed in the 1990s as a way to represent meaning in computational linguistic applications. The procedure closely followed human understanding of natural language, which was largely compositional. Native speakers intuitively assemble sentences out of the words they know; these meaning interact in established way, and then the sentence means whatever the speaker or writer needs to express.

It was informed by two intuitions about meaning that were new and not very popular. First, there was Raskin's idea that language meaning was organized into scripts/frames rather than being confined to separate words. The same idea was being developed by several people at the time and was probably initiated much earlier in psychology (see [1]-[3]). It was applied and somewhat developed for the first linguistic theory of humor [4].

Second, in the context of the briefly resurrected machine translation, Nirenburg and Raskin, in the 1980s, put forward the idea of interlingua as a mediating system between a source language and a target language. This appeared to be much smarter than the transfer systems translating only between a pair of languages in one direction at a time. It made perfect sense for the interlingua to be usable in any pair of languages, and for that, it had to be a semantic representation.

Machine translation was the process of translating text in a natural (source) language into its meaning representation and then from that representation to another natural (target) language. Immediately, the major issue hampering the development of linguistic semantic for centuries raised its head: What was the medium of the semantic interlingua? Increasingly, various private and government groups were talking about ontologies, meaning primarily inventories of terms. The purpose of these systems was largely the standardization of terminology, and most, if not all items, were nominals. The funding came first from NSF and later, massively, from NSA. [5] somewhat timidly, referred to our interlingua as an ontology, and it took.

Ontological Semantic ontology, in its numerous incarnations, contains properties and concepts, and in that it is somewhat similar to other ontologies in the semantic web.

 ISBN: 978-1-61208-705-4

However, as any ontologist would know, ontologies are not compared by the number of concepts or even the number of properties that they support, it is how the concepts and properties are used for the reasoning capabilities that should be compared. In this sense, Ontological Semantic is very different from a well-known CyC or various ontologies of description logics.

An initial massive acquisition of ontology took place at the Computing Research Laboratory at New Mexico State University which was inherited by Nirenburg from Yorick Wilks, a somewhat apprehensive friend of Ontological Semantics, and Raskin became a regular PI-level consultant there. The collaboration took place in 1994-2002 and was supported by a variety of grants, mostly from NSA.

By 2001, the rich experience, historical but mostly ideational, was summarized in the monograph distributed generously online and published later [6]. The book's title, "Ontological Semantics," gave the name to the 1990s approach and resources, though sometimes they have been referred to, reviewed, followed, and criticized as "Mikrokosmos," which was the title of a central grant and conformed to the CRL tradition of Greek names for computers and other items—hence, the k's in the spelling.

Figure 1 shows the architecture of a later developed Ontological Semantic Technology (OST) system, as it interprets every sentence ontologically. The ontology contains concepts, which are mostly events and objects, linked with named properties, all connected hierarchically on the ISA property. The ontology is language independent, which means that it is the same for all languages, the concepts are labeled in English; the labels are not English words in that they are not polysemous, have no synonyms, and are not understood by the computer; they just name the unique nodes (which also have IDs) and their links and are convenient to use for the presumably English-reading acquirers. The specific language information is acquired in the language-specific lexicon supported by morphological and syntactic knowledge (phonology is added for specific applications).

While the ontology is language-independent, the lexicons are language-specific. It is worth noting, however, that while word-for-word translation is somewhat useless, an empirical evidence suggests that sense-for-sense translation makes the development of another lexicon an easy 6-person-month project by a bilingual undergraduate.

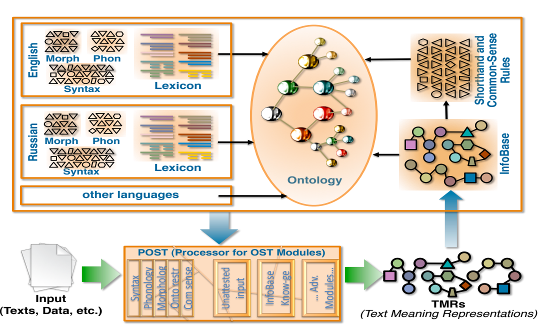

Figure 1: OST diagram as reported in most post 2012 OST papers

The resources are acquired semi-automatically supported by an online resource that shapes and automatically fills out the entries. The human user is asked only to exercise judgment about meaning. In ontological acquisition, the human determines, very importantly, where to find it in the ontology, possibly under a different label, and in case of failure, where to add it as "child," or leaf, of an existing concept. The difficult decisions, such as opening high-level leaves or making changes and pursuing them consistently, are made by Master Ontologists, of whom the current co-authors are two (of possible three or four altogether).

An obvious argument here is if it was easily done, why is there only several Master Ontologists. The answer is opportunistic and fashion-related: with the machine learning producing promising results, the race for using ML methods attracted and pulled in the majority of researchers whose careers required publications (with popular methods). It is only now that the field slowly realized that deep learning and machine learning are not an answer to all questions, and going back to and incorporating the so-called-first-wave of AI might be a good idea [7-8].

Various versions of Ontological Semantics exist today, some can be accessed online. Figure 2 illustrates a few top ontology layers of the current version of Purdue online resource, available at engineering.purdue.edu/~ost, implemented by several of Rayz's undergraduate students. Each node can be clicked to show the sub-hierarchy, if any, of which it is on top.

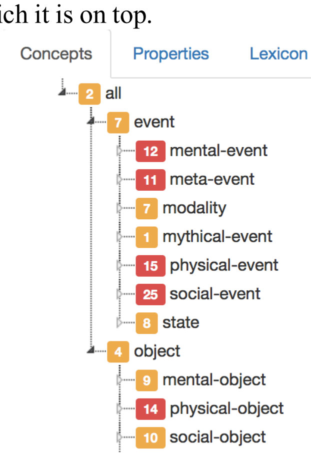

Figure 2: Upper level hierarchy of ontological concepts, Purdue OST webtool

The lexicon is acquired by native speakers, and the training is limited to a session or two for the basic acquisition. Most nouns and verbs are anchored in a concept, and the concept frame prompts the acquirer to the related words in their slots. Thus, the verb *enter* will want to identify a typical agent, theme, instrument, etc. A typical adjective will be anchored in a property, and those are defined in terms of domains and ranges, that is, the object or events they define and what values they receive. All that information is contained in the resource, and the acquirer is prompted for it automatically.

A simple English sentence *John is driving to the store for groceries* will be transformed onto this simplified format (Figure 3), where all items are ontological:

go
agent John
instrument car
Destination store
item-for sale groceries

Figure 3: Conceptual representation of a TMR

This schema is a Text Meaning Representation (TMR), and the purpose of obtaining TMRs for all sentences was basically achieved in principle and for limited applications by the late 1990s. In 1997, NSA published an RFP that called upon all applicants to follow the ontological semantic format. They all said they would but…

III. VECTORS AS MEANING

By the time Raskin distributed locally the mimeographed edition of his "Concise History of Linguistic Semantic" in 1973 [9], he included the one then known example of a statistical approach to linguistics. A. Shaykevich [10] obtained a copy of the published concordance of Shakespeare's complete works where every line was numbered throughout his legacy and every adjective was listed with the line numbers where it occurred, He then calculated pairs of adjectives cooccurring much more than statistically; he assigned them links that were reversely proportionate to these ratios, so that, on a graph, they were closer to each other. The result was spectacular: they were whole areas of semantically linked adjectives. Encouraged, he took the show on the road, and Raskin was one of many session chairs who tried to get him to finish the paper on time. That was a rare bird then: linguists pretty much ignored the statistical methods and did not accept their papers at linguistic conferences. So they developed outside of and aggressively without linguistics, done not by trained linguists but by statisticians, programmers, "techies."

By the early 1990s, at least partially energized by Rumelhart's ideas of the 1980s, the statisticians broke the dam, and within a decade, they came to completely dominate the field of natural language processing. Originally presented as the alternative to rule-based approaches, solidly rooted in linguistics, the first wave of machine learning, rapidly growing through computer science and engineering graduate programs, came up with packs of inventive algorithms that quickly and reliably identified outliers in a large sample of data, and they did it without any effort on linguistic resources like lexicons.

The approach first claimed boldly that it was not interested and did not aspire to represent what the text meant. This was no longer the task or purpose of natural language processing. Starting with machine learning, it evolved into neural networks, word embeddings, and deep learning. The purpose of each undertaking is to obtain vectors that are somehow presented as meanings.

We fully realize that the goals of old-school computational semantics is very different from all these approaches, the technical and formal virtuosity of their best work is admirable, and finding the attempts to semanticalize their scope somewhat hopeful. Thus, the word2vec [11, 12], being the first of word embedding papers, have rediscovered Firth, the last structuralist to try and define semantics "distributionally," that is non-semantically. Famously at the time but then forgotten until now, Firth [13] described the ability of the word *night* to be combined with *dark* as part of its meaning, and, of course, vice versa. This can be extended to saying that the meaning of dark is its ability to be combined with all the words it is combined with. This will blur several different meanings of dark together (*dark mood, dark comedy*) and describe it non-substantively. In fact—and it was not seen then—it is not significantly different from Wittgenstein's defeatist statement that language is use ([14]—see also Section V below). Interestingly, within several years of word embedding promising results, papers staring to appear showing that word embeddings do not correspond to human knowledge and reasoning (see papers in Cognitive Science conferences starting with 2016). Ideally, we wonder if phenomena discovered and defined substantively ontologically can then be identified on large data with the best of automatic methods.

IV. EXTENSION I: FROM APPLIED TO THEORETICAL SEMANTICS

Let us forget about vectors substituting for meaning and return to representing what sentences in natural language mean for the users.

The TMR on Figure 3 represents, like any one TMR, a number of sentences that are paraphrases (shown below), and those are numeral in any natural language (in the 1960s, Mel'cuk calculated that a regular Russian sentence stating that it was hard for Smith to translate the text because there were many technical terms in it could be paraphrased in over a million ways). In this case, the paraphrases include sentences like (i-iii) and many more.

(i) John is driving to the grocery store
(ii) John is driving to a store to buy groceries
(iii) John is driving for groceries

There are many other things that are closely related and are assumed by the speakers: that groceries are in a store, that the store must be open in order for John to achieve the goal, that the store is reasonably close, that John will have to park the car before going to the store, etc. These are often referred to as inferences, and they are all part of what humans understand when they process the largely compositional meaning of a sentence consisting of certain words. The same idea can and often is expressed differently, while keeping or perhaps slightly modifying the inferences.

What is clear is that explicit semantics assigning a TMR to a sentence fails to express clearly and accurately what a sentence actually means because meaning is largely implicit. The recall experiments of the 1970s [15] demonstrate that, a week after being exposed to a verbal event, the subjects remembered the exact wording differently while all reproducing the gist of the event accurately without differentiating much among the inferences.

This is primarily how and why it dawned on us, in the late 2000s, that Ontological Semantics is not limited to computer applications and to computational semantics in them: it is, in fact, the theoretical basis of semantics in general—it is the only reliable way to handle meaning. We call the new applied approach Ontological Semantic Technology (OST: [16]-[18]), developed and extended for a while to a couple of high-tech start-ups. We were still developing semantic applications and resources for them but, theoretically, we were enlarging the scope to scripts and other complex semantic phenomena, often deliberately limiting the scope of the approach to specific applications, past and current. In both theoretical and applied undertakings, we are often extending the scope to verbal humor, a structured form of discourse, leading to promising scalable applications in computational humor [19].

V. Meaning Substance

One persistent problem with semantics, both as a linguistic subdiscipline and a branch of philosophy and mathematics but especially the first one, is the elusive nature of its substance. Meaning is real but what is it? It is something humans know, share, and convey but the messy and informal mechanism of paraphrase is the only manifestation of that understanding.

Bloomfield [20] is often referred to as having excluded semantics from linguistics even though his monograph has a chapter entitled "Meaning." He was the one who claimed counterintuitively that in order to understand the meaning of the English word *pie*, not only to we have to know all of its ingredients and how they were baked together but also the state of each of its molecules at all times. He saw the meaning of Jill's request to Jack to get her an apple from a tree as a replacement for the extralinguistic substance of stimulus and response: Jill sees an apple, she is hungry but instead of responding to that directly by picking it she replaces that with a linguistic stimulus.

Yngve [21] reiterated his similar desire to replace semantic substance with that of scientific observation and, capitalizing on a good friendship over decades, invited Raskin to join him there. Raskin had to decline because it was clear to him that no matter how well his behavior is observed and recorded, there is no way to know what is being said or written: the substance of observation is much too coarse.

In semiotics [22], semantics was introduced as a component where the items were interpreted. The examples included interpreting some variables were interpreted as named contacts. This is how semantics is used formally in logic except that it is not used there much. The substance of the interpretations never needed elucidation, so there was none.

At the very late beginning of linguistic semantics in the 1860s, the discipline happily and innocently substituted purview for substance. It included multiple facts, from an assortment of languages, about words changing their meanings historically. Those meanings, both past and present, were outlined very approximately, usually with one reference to a class term, such as clothes for Latin *vestis*, later English *vest,* a garment.

When, next, semantics started studying the nature of meaning, the purview got limited to Frege's distinction between, roughly, meaning and reference, neither of which were well-defined. The most significant effort toward discovering the substance of semantics had to do with the field's 20-30-year romance with semantic features within the componential analysis approach of the 1930-1950s. The exciting mathematical idea that 21 binary features can describe over a million items has rather quickly lost its attraction when it turned out that, in semantic reality, many features have a very limited scope, such as 'never-married' describes the meaning of *bachelor* and its awkward female counterparts *spinster* and *old maid*. Also confusing to critics was the labeling issue: because they were named with English words, Kats and Fodor's [23] semantic markers were dismissed by Lewis [24] as Markerese for no apparent reason.

Within sentential semantics, formalism, rules replaced any serious concern with substance, and foundational considerations have been abandoned or relayed to mathematical logic, which is where the formalism originated from. Somewhat in that tradition, Barwise and Perry [25] attempted to replace the two truth values only as the range of extension of a proposition with facts, that is what the proposition was about. This would have revived perhaps the intension/extension debate, with benefits to various semantics, but the authors could not withstand the ferocious attacks from philosophers about how fact was defined.

Yet meaning is definitely about something other than the words used to express it. Sentences state something about items mentioned in them. Marked in a natural language, this content is independent from it. Living in bilingual environments, we are both perfectly aware of that, and the fact that everything is double-coded barely interferes with our non-linguistic perception of reality. When interpreting between languages, we do not replace words with their translation but rather reconstruct the reality from the sentence in the source language and then express it in the target language (in the process, incidentally, some words are replaced by their translations).

When one learns about a new piece, area, or domain of reality one needs to identify the major agents and events they participate in. We have gone over that ontologically whenever our ontology needed an extension. A major text in the field provided an index where we started. We extended the then current ontology into banking and into information security, and it took under 6 person/months of doctoral labor at the cost of under $20K. The users were implementing a very specific well-defined job, and the domain reality was definitely not a linguistic object (incidentally, both acquirers were, of course, English speakers but with different native tongues; one of them later singlehandedly extended the single ontological node, FEELING, into all other feelings from psychology, and in much more detail, into the field of humor research—see [26]. Similarly, an ontology was expanded to a domain of phishing detection, with a similar effort, producing successful results in phishing detection [27].

VI. EXTENSION 2: ONTOLOGY AS HUMAN REALITY

The disparate considerations in the previous section seem to indicate that people can think of various pieces of reality as separate from the languages in which they describe them. We propose then to see the ontology as the extralinguistic structure of reality and the medium of semantic substance. This is already how we have used it prior to this claim. What is different now is that we present our ontology as the description and representation of human reality in the computer. Quite simply, what the computer knows about the world depends on its ontology, both explicit and implicit, and we are interested in doing that explicitly. The concept and role of ontology is thus extended from an application-focused tool that works well within it to the theoretical basis of all meaning studies and, finally, to the structure of the world in the computer.

This is not very easy to understand because people easily confuse the accumulation of large data in the computer storage with what the computer knows. IBM Watson is a good example. The promotion materials and journalists easily describe it as knowing an awful lot but the serious founders carefully explained in the NPR Nova program before the introduction of the system on "Jeopardy" back in 2011 that it was devoid of any intelligence. What Watson could do in 2011, with amazing technical speed and dexterity, was accessing everything from its enormous storage which has the same words as the query and manufacturing a response on this basis. It had no idea what the question was about nor what its memorized quotes say.

A more recent example is Amazon's Alexa. As of March 2019, Alexa is perfectly capable of telling a user about the weather (in whatever scale the user specified in the setting), it can also convert from one temperature scale to another. What it cannot do, however, is tell the weather in the scale that a user asked in his/her question. In other words, information retrieval works very well, but any additional manipulation of the retrieved information presents a challenge.

Ontological competence underlies inferences. The more intelligent the person is the more inferences are available to it. Unlike the computational inference engines of yore (read: 1980s), people do not generate inferences combinatorially: besides being guided by ontological links, they have an ability to cut through to the relevant ones only, and we need to understand this capacity better and to emulate it in the computer, and statistics will not help us here.

Alexa, Siri, and other personal assistant software are usually listed in the media as successes in artificial intelligence. It can control many smart home devices, including a smart thermostat when explicitly asked to do so. However, it cannot answer the question, *Why am I cold*?—with something like, *Do you want me to raise the setting inside*? At the same time, Alexa constantly evolves: new pieces of knowledge are added, some of them bizarre. It can translate some days of the week into Russian, but not all. It knows who Anna Akhmatova is, likely surprising its customers including most in post-Soviet Russia, who are much better familiar with Beyoncé. The question is, how are these disparate pieces are selected to be added? Is there a selection process? How does one determine what is useful and what is not?

With ontology, the process of extension and filling the gaps is guided, and the ontology is improved and corrected with every new text that is processed and new TMRs generated. We achieve here a new theoretical level of completeness: a theory is complete if it has handled everything well so far, and it is indefinitely extendable. Much of human knowledge is infinite, and its adequacy is temporary.

The improvement record in Ontological Semantic Technology is constant, systematic, and reducing in volume. A small group of experts has to be maintained to take care of it, though it will be increasingly automated. One may try and argue that trying to fill the gaps ontologically is as haphazard as what is being done with Alexa. The difference is the links in the ontology, which lead to all possible inferences directly or transitively.

The most important difference between the ontology as we do it and any other formalism is isomorphism between every single TMR and a reality fact (pace Barwise and Perry) that takes place or may take place or not take place in the world, including the extensions to myths, fantasies, and counterfactuals. The ontological items, when done right, relate to each other exactly as in real life: whatever happens (or does not) in the world is reflected ontologically. And the ontology as the computer knowledge of the world is constantly checked, corrected, and upgraded in computational linguistic applications on demand.

VII. CONCLUSION

We have guided the reader briefly through a different and difficult terrain. Contrary to the dominant view that semantics is unknowable and the knowledge of meaning should be replaced by machine learning methods and their extensions, we propose a view that it is accessible. This view provides explanation for every decision it makes, and produces reasoning scenarios that are needed for the "third wave" of AI. Then, we take you to a *terra incognita* of the semantic substance: what it is that the meaning is interpreted and presented as. The answer is, of course, the conceptual hierarchy with many properties, the comprehensive ontology. After that, it is almost easy to claim that this ontology represents our knowledge of the world. And because it is a formal object it can be introduced into the computer so that it also knows the world.

REFERENCES

[1] F. C. Bartlett, Remembering. Cambridge: Cambridge University Press, 1932. Reprinted by Cambridge University Press, 1977.

[2] M. Minsky, A framework for representing knowledge, in The Psychology of Computer Vision, P. H. Winston, Ed. New York: McGraw Hill, pp. 211-277, 1975.

[3] R. C. Schank and R. Abelson, Scripts, Plans, Goals, and Understanding. New York: Wiley, 1977.

[4] V. Raskin. Semantic Mechanisms of Humor. Dotdrecht-Boston-Lancaster: D. Reidel.

[5] L. Carlson and S. Nirenburg, World Modeling for NLP. Technical Report CMU-CMT-90-121, Center for Machine Translation, Carnegie Mellon University, Pittsburgh, PA, 1988. A short version appeared in Proceedings of the Third Conference on Applied Natural Language Processing, Trento, Italy. April.

[6] S. Nirenburg and V. Raskin, Ontological Semantics. Cambridge, MA: MIT Press, 2004.

[7] A. Prabhakar, Powerful but Limited: A DARPA Perspective on AI, https://sites.nationalacademies.org/cs/groups/pgasite/documents/webpage/pga_177035.pdf.

[8] J. Launchbury, A DARPA Perspective on Artificial Intelligence, Machine Learning, 2017.

[9] V. Raskin, A Concise History of Linguistic Semantics. Tel Aviv: Tel Aviv University, mimeographed, 1973. Reprinted at Purdue University, 3rd ed., 1983.

[10] A. Ya. Shaykevich, Vydelenie klassov slov i paradigm posredstvom distributivno-statisticheskogo metoda (na materale komediy Shekspira) /Identification of classes of words and of paradigms with the distributive-statistical method (on the material of Shakespeare's comedies), Prikladnaya Lingvistika, Vol. 18, pp. 96-134, Moscow, 1976.

[11] T. Mikolov, K. Chen, G. Corrado, and J. Dean, Efficient Estimation of Word Represenetation in Vector Space, ICLR Workshop, 2013.

[12] T. Mikolov, I. Sutskever, K. Chen, G. Carrado and J. Dean, Distributed Representations of Words and Phrases and their Compositionality, NIPS 2013.

[13] J. R. Firth, Modes of meaning, in his Papers in Linguistics, London: Oxford University Press, 1964.

[14] L. Wittgenstein, Philosophical Investigations, Oxford: Blackwell, 1953.

[15] W. Chafe, Language and memory, Language 49:2, pp. 261-281, 1973.

[16] J. M. Taylor, C. F. Hempelmann, and V. Raskin. On an automatic acquisition toolbox for ontologies and lexicons in ontological semantics, Proceedings of the International Conference on Artificial Intelligences, Las Vegas, NE, pp. 863-869, 2010.

[17] J. M. Taylor, V. Raskin, and C. F. Hempelmann, From disambiguation failures to common-sense knowledge acquistion: A day in the life of an ontological semantic system, Proceedings of Web Intelligence Conference, Lyon, France, 2011.

[18] J. M. Taylor, V. Raskin, and C. F. Hempelmann, Towards computational guessing of unknown word meanings: The ontological semantic approach, Proceedings of the Cognitive Science Conference, Boston, MS, 2011.

[19] V. Raskin, C. F. Hempelmann, and J. M. Taylor, How to understand and assess a theory: The evolution of the SSTH into the GTVH and now into the OSTH, Journal of Literary theory, 2009 (as marked but came out in 2010).

[20] L. Bloomfield, Language. New York: Holt, 1933.

[21] V. Yngve and Z. Wasik, Hard-Science Linguistics, New York: Continuum, 2006.

[22] Ch. Morris, Signs, Language, and Behavior. Chicago: University of Chicago Press, 1946.

[23] J. J. Katz and J. A. Fodor, The structure of a semantic theory, Language, 39: 1, pp. 170-210, 1963. Reprinted in The Structure of Language: Readings in the Philosophy of Language, J. A. Fodor and J. J. Katz, Eds. Englrwood Cliffs, NJ: Prentice-Hall, pp. 479-518, 1964.

[24] D. Lewis, General semantics, in Semantics of Natural Language, D. Davidson and G. Harman, Eds. Dordrecht-Boston, D. Reidel, pp. 169-218, 1973.

[25] J. Barwise and J. Perry, Situations and Attitudes. Cambridge, MA: MIT Press, 1983.

[26] K. E. Triezenberg, The Ontology of Humor, Unpublished Ph.D. Dissertation, Program in Linguistics, Purdue University, West Lafayette, IN, 2006.

[27] G. Park and J. Rayz, Ontological Detection of Phishign Emails, IEEE-SMC 2018.

 ISBN: 978-1-61208-705-4