Computer Science Journal of Moldova, vol.**24**, no.3(72), 2016

Natural Language Processing versus Logic. Pros and cons on the dispute whether logic is useful in the computational interpretation of language

Dan Cristea

Abstract

In this essay I express some personal opinions regarding the influence that logic has on modern approaches to process natural language in artificial systems. I start by presenting some successful linguistic formalisms that were originated in logic, arguing why logic is important in conveying the meaning of language expression. Then, I counterbalance the argumentation with a number of examples where logic is impuissant to mirror language usage, finally supporting a rather temperate opinion about the usefulness of logic to formalise low level linguistic processes and about the limits of language formalisation.

Keywords: natural language processing (NLP), logic theories in NLP, statistical approaches, symbolic versus statistical approaches in NLP.

1 Introduction

It seems that we, human beings, motored by the need to understand the reality among us (a condition for survival), make use of shallower or deeper cognitive processes in our efforts to assign meanings to messages we receive through language. This cognitive behaviour resembles, in some cases, logical formalisms (what is logic if not an expression of thinking?). This is just as we have in our brains a symbolic machinery capable to help us make inferences and offer solutions to complex puzzles that the language encodes. In other cases, however, logic is of no use: language manifests a totally weird behaviour.

The need for symbolic approaches, as opposed to statistical ones, is in itself an expression of the feelings researchers have that language can be expressed as a system of rules, which would make its messages to be "computable". This paradigm is similar to the one of mathematical logic, since there too, based on a very clear notation of a number of basic ingredients and of defined ways in which they can be grouped together, truth values of sentences (i.e. syntactical constructions) can be deduced.

The debate here, is not if logical, i.e. symbolic, formalisms are useful, but to what extend can they be applied to explain a range of linguistic phenomena. Then, for the whole rest of linguistic phenomena for which logic fails to offer support, ought we instead resort to statistical or neural-based solutions? Or part of our manifestations triggered by language escapes from any formalisation, being it based on logic, statistics or neural grounds? And finally, what is the range of practical applications of language which put at their base logical solutions?

In this essay I express some personal opinions regarding the influence that logic has on modern approaches to processing natural language.

2 When do we need logic to decipher language?

Language ought to be logical, or, else, the communication based on it would be impossible (imagine that the inscription of messages would be made in a randomizing system of signs...). In most of cases, people are able to transmit their intentions correctly to the intended receivers. So, when we say *two horses* we mean something of the kind: $\exists\, S = \{x|horse(x)\}^{\wedge} card(S) = 2$, where $card(S)$ means the cardinal of the set S, $horse(x)$ is a qualifying function asserting that x, its argument, has a semantic property that is shared in the common knowledge of the living inhabitants of this world, that this property is currently denominated in English by the word *horse*, and that the pragmatic context in which the expression is uttered clearly separates the different meanings that this English word may encode. In the same time, uttering this noun phrase, we are also aware that the listener possesses an equivalent decoding mechanism that enables her to coagulate an equivalent meaning. So, it seems that in extremely many situations of the real life, when we use language, we implicitly resort to mathematical logic to express our messages. In fact, although we don't really do that, it is just like doing that, i.e. just like somebody above us, listening to what we are saying, quickly encodes our saying in a logical expression that could unambiguously be "read" or decoded by our partner in the conversation.

The challenge to describe language in a logical system of notation that would allow non-ambiguous representations of its lexical elements and coherent composition/decomposition has preoccupied modern linguistics for a long time already. In this section I will make a very quick survey of only some of these approaches.

2.1 Generative Lexicon and Qualia Structures

This kind of attitude towards language brings forward Generative Lexicon (GL) [6], a theory that intends to build lexical and semantic resources capable of expressing in computational terms (which is another name for logic) the rich lexical variety of the language (any language, in principle), including its capacity to combine meanings of lexical items through grammar and, to a certain extend, through pragmatics. GL tries to describe the semantic flexibility shown by words in combination with others. To account for diverse interpretations that words can display when placed in combinations with others, GL associates a hidden event in the lexicon description of nouns, adjectives and adverbs. Originated in the Aristotelian concept of *aitia* (explanation), with reinterpretations added by Moravcsik [4], further developments of GL [5, 7] introduce Qualia structures, describing roles according to which the meaning of words can be decomposed on four different coordinates:

- *Formal* (F): encoding taxonomic information about the lexical item (the *is-a* relation); □
- *Constitutive* (C): encoding information on the parts and constitution of an object (*part-of* or *made-of* relation); □
- *Telic* (T): encoding information on purpose and function (the *used-for* or *functions-as* relation); □
- *Agentive* (A): encoding information about the origin of the object (the *created-by* relation). □

Qualia are formally represented as typed feature structures. For instance, the one in Fig. 1 can account for combinations such as: *large car* (F = vehicle), *broken car* (C = motor), *speedy car* (T = drive), *Italian car* (A = made in Italy). Also, adding arguments to roles, Qualia structures can deal with metonymy, as in: *the car from behind honked*: T = drive (human, vehicle). Such lexicon representations actually encode in a logical form part of an ontology of lexicalised concepts, out of which "understanding" can be computed.

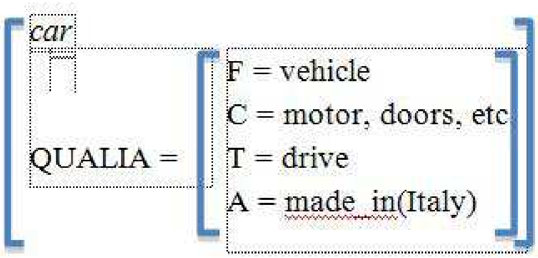

Figure 1. A Qualia structure for the lexical concept car

Qualia structures are recognised to have various shortcomings. Pustejosky and Jezek [7], for instance, recognise the rather limited "ability [of the formalism] to take on an indefinite variety of possible senses depending on the other words they combine with". He gives the example of the verb *like* and wonders whether it has two different meanings in "He likes my sister" and "He likes vanilla ice cream", and if so, how is this difference to be represented in decompositional terms?

2.2 Discourse coherence

But logic is needed at upper levels of language interpretation as well. One example is the need to consolidate meanings of sequences of utterances in discourse. Compare, for instance:

Maria dropped the egg from her hand. (1.1)
She cleaned the floor. (1.2)

with:

Maria dropped the feather from her hand. (2.1)
She cleaned the floor. (2.2)

While sequence (1) is perfectly coherent, sequence (2) apparently has no meaning, although the utterances of each sequence convey unambiguous meanings. It is clear that an inferential chain of deductions, triggered by common sense (or ontological) knowledge, link the two utterances in (1), as opposed to (2), where the connection is much harder to establish. The meaning in (1) is built out of a reasoning sequence showing a temporal occurrence of events that could be schematized as follows:

(1.1) $\Rightarrow$ drops(AG:Maria, OB:egg) $\Rightarrow$ falls(REC:egg) $\Rightarrow$
touches(AG:egg, OB:X) (3)
Qualia.Formal(egg) = container; Qualia.Constitutive(egg) =
{eggshell(fragile),liquid} (4)
(3), (4) $\Rightarrow$ breaks(REC:eggshell) $\Rightarrow$ leaks(REC:liquid, ON:X) $\Rightarrow$
perceives(AG:Maria, OB:dirty(X)) (5)
(1.2) $\Rightarrow$ cleans(AG:Maria, OB:floor) $\Rightarrow$
perceives(AG:Maria, OB:dirty(floor)) (6)

The equality of (5) and (6), when X equals floor, closes the inference chain, proving the high degree of coherence of the sequence (1). A longer inference chain (if any), implying volitional searches in a space of possibilities fuelled by imagination, as opposed to the first case, in which the inferences are common-sense, natural, spontaneous, could, in the mind of an intrigued reader, possibly link the utterances in (2), thus showing a much lower degree of coherence.

2.3 Textual Entailment

In the fight to decipher the meaning expressed in language, two contrary phenomena have to be faced: variability and ambiguity. Variability of language means that the same meaning can be verbalized in different surface forms. Ambiguity means that one surface form can be interpreted as having different meanings. A number of NLP applications that deal with the variability of language trying to reduce the distance between form and meaning are: Information Retrieval (IE), Textual Entailment (TE) and Question Answering.

In TE, it is said that **text *t* entails hypothesis *h* ($t \Rightarrow h$) if humans reading *t* will infer that *h* is most likely true**. So, textual entailment is a directional relation between two texts. In practical applications of TE (including competitions[1]) t could be complemented with external knowledge in order for h to be entailed, but h cannot be entailed only by the knowledge itself (for instance, by searching on the web).

Here is an example of a true entailment (from RTE data):

t: *... a shootout at the Guadalajara airport in May, 1993, that killed Cardinal Juan Jesus Posadas Ocampo and six others.*
h: *Cardinal Juan Jesus Posadas Ocampo died in 1993.* (7)

and of a false one (same source):

t: *Regan attended a ceremony in Washington to commemorate the landings in Normandy.*
h: *Washington is located in Normandy.* (8)

One of the methods used to measure the similarity between t and h does syntactic matching or transformations at the syntactic level. To

[1]For instance, the EU FP-6 Funded PASCAL Network of Excellence 2004-7: Recognizing Textual Entailment (RTE) Challenges.

see the complexity of such an attempt, I will examine in some detail an example. Suppose t is:

Philanthropic Golding Inc. came into existence in January 2004. (9.1)
One year after its foundation the company declared bankruptcy. (9.2)

and h:

Philanthropic Golding Inc. bankrupted in January 2005. (10)

One way to check the validity of such an entailment, is to launch a pipeline of processes, at the end of which the sentences of both t and h are expressed in a symbolic form that allows close comparison. Applied to (9.1) the pipeline produces the following successive results (simplified)[2]:

Step1: tokenisation (not shown), part-of-speech tagging (not shown), chunking noun phrases and clashing multi-word expressions.
<NP id="n1">*Philanthropic Golding Inc.*</NP>
<MWE id="m1">*came into existence*</MWE>

Step 2: recognition of entity mentions, of time expressions and resolution of anaphora.
<COREF-LIST id="ent1" TYPE = "ENTITY" REF-LIST="n1" />[3]
<TIMEX3 tid="t1" type="DATE" value="2004-01">*January 2004*
</TIMEX3>

Step 3: functional dependency parsing; in Figure 2 we show a Universal Dependency (UD) coding [3].

Step 4: time analysis, in which EVENT and TLINK elements, formalizing the events and their temporal relations, are generated[4].
<EVENT eid="ev1" VB="m1" AG="ent1"/>
<TLINK eventID="ev1" relatedToTime="t1" relType="BEGINS"/>

[2]Here we use an XML coding, but a representation that uses RTF tuples or another notation convension can also be employed.

[3]A COREF-LIST with only one member signals the first mention of an ENTITY or EVENT, according to TYPE.

[4]To simplify notations, MAKEINSTANCE and SIGNAL elements are ignored and EVENT elements are complemented with roles.

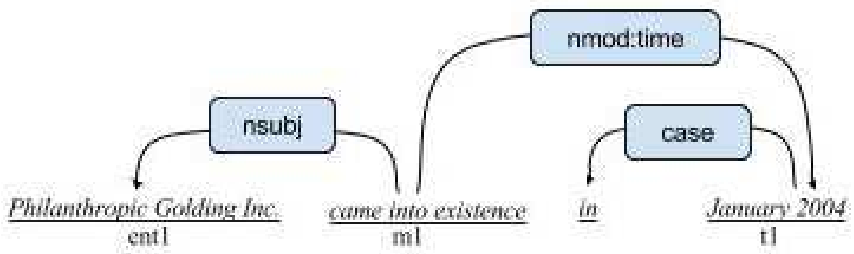

Figure 2. UD parsing of sentence (9.1)

Step 5: generation of equivalent structures (transformations) and the application of the closure tool, which computes the transitive closure of temporal relations; the transformation here concerns the equivalence of the expressions: *X comes into existence* and *UNKNOWN founds X*. This rule triggers the element:

<EVENT eid=”ev2” VB=”found” AG=”UNKNOWN” OB=”ent1”/>

and its correspondent time link:

<TLINK eventID=”ev2” relatedToTime=”t1” relType=”BEGINS”/>

And now the pipeline applied to (9.2):

Step 1:
<NP id=”n3”><NP id=”n2”>*its*</NP>*foundation*</NP>
<NP id=”n4”>*the company*</NP>
<NP id=”n5”>*bankruptcy*</NP>

Step 2:
<TIMEX3 tid=”t2” type=”DURATION” value=”P1Y”>*one year*
</TIMEX3>
<COREF-LIST id=”ent1” TYPE=”ENTITY” REF-LIST=”n1 n2 n4” />
<COREF-LIST id=”eve1” TYPE=”EVENT” REF-LIST=”ev1 ev2 n3” />
<COREF-LIST id=”eve2” TYPE=”EVENT” REF-LIST=”n5” />

Step 3:

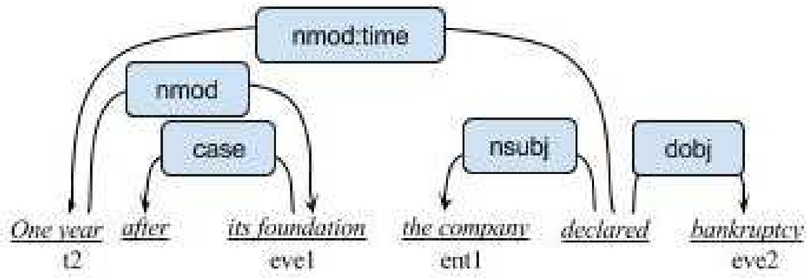

Figure 3. UD parsing of sentence (9.2)

Step 4:
<EVENT eid=”e3” POS=”VERB” CLASS=”REPORTING” AG=”ent1” OB=”eve2”>*declared*</EVENT>
<EVENT eid=”e4” POS=”NOUN” CLASS=”OCCURRENCE” AG= ”ent1”> *bankruptcy*</EVENT>
<TLINK eventID=”e3” relatedToTime=”t2” relType=”AFTER”/>
<SLINK eventID=”e3” subordinatedEvent=”e4” relType= ”FACTIVE”/>

Step 5: the closure tool produces:
<TIMEX3 tid=”t4” type=”DATE” value=”2005-01-xx” />
<TLINK eventID=”e4” relatedToTime=”t4” relType= ”DURING”/> (11)

A similar processing pipeline applied to (10) should yield:
<EVENT eid=”e5” POS=”VERB” AG=”ent1”>*bankrupted* </EVENT>
<TIMEX3 tid=”t5” type=”DATE” value=”2005-01-xx”> *January 2005* </TIMEX3>
<TLINK eventID=”e5” relatedToTime=”t5” relType= ”DURING”/> (12)

And the equivalence of (11) and (12) proves the entailment.

2.4 Other NLP formalisms rooted on logic

Prolog, the programming language of logic, has inspired much work on NLP. In syntax, this means to express a grammar as a set of statements in a logic formalism (e.g. Horn clauses), and to use a theorem prover (e.g. resolution) in order to parse or generate sentences. Recently used in information extraction, SHERLOCK [8] is a system able to learn Horn clauses in a large-scale, domain independent manner, from Web texts. The learned rules can then be used to fuel a first-order reasoning system, as HOLMES, described by Schoenmackers *et al.* [9], which infers answers from tuples.

3 And when logic is of no use?

Languages have specific ways to express linguistic phenomena. Some of them seem to escape any logical explanations.

3.1 Double negation

In propositional logic the double negation is equivalent to an affirmation. However, applied to language, this rule does not always hold. In connection to this phenomenon, Falaus [2] inventories two main types of languages. In Double Negation languages, among which standard varieties of Germanic and Scandinavian, two negative elements cancel each other out resulting in a positive reading, as in (13) below:

Paul didn't see nobody. = *Paul saw somebody.* (13)

However, in Negative Concord languages, among which Romanian and Italian, multiple occurrences of negation are interpreted as one semantic negation, as in (14):

Paul n-a văzut pe nimeni. = *Paul didn't see anybody.* (14)

The sentence can be paraphrased as "It is not the case that there is an individual x, such that Paul saw x."

3.2 Linear position

Romanian is known to be ambivalent with respect to the position of quality adjectives around the nouns they modify: they may occur pre-nominally as well as post-nominally. Cornilescu [1] notices that certain associations of noun+adjective versus adjective+noun makes a difference of interpretation, as here:

femeia singură => the woman alone
singura femeie => the only woman (15)

I believe that the following examples display similar behaviour. Suppose somebody has two cars, one bought some time ago and one recently bought, and the one recently bought belongs to an old brand while the one bough in the past belongs to a newer brand. Then:

maşina lui cea veche refers to *his old brand car*, while
vechea lui maşină refers to *the car owned by him for a long time* (16)

However, in the following associations the sense does not change:

domnişoara frumoasă, frumoasa domnişoară => the beautiful young lady
cartea interesantă, interesanta carte => the interesting book (17)

Also the positional ambivalence does not apply to any adjective. Certain modifiers make sense only when situated in the pre-position with respect to the modified noun. For example, *biet* (*poor*, *pitiful*) is not accepted unless it precedes the noun: *biet om* (*poor man*), but not: *om biet*. Vulchanova [10] explains this for Balkan languages: these associations seem to contradict the usual intersection-based composition. In general, if X is an adjective and Y – a noun, then XY (or YX) means the set of objects Y that have the property X, or the intersection between the set of objects having the property X and the set of objects Y. As such, *poor men* should be taken as the intersection between the set of things which are poor and the set of men, but *bieţii oameni* means something different than the subset of the set of men which are poor, it means a subset of the set of men which are in a pitiful/miserable state.

3.3 Contexts and the mist of pragmatics

It is a truism that the context determines the meaning of words, and the previous section showed some examples. By "context" here I mean both textual (i.e. positional) and not textual (for instance, temporal).

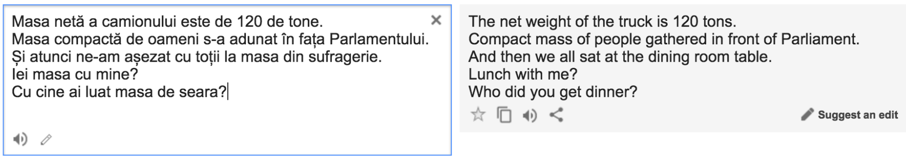

Figure 4. Google Translate solutions for difference occurrences of the Romanian word "masa"

The left side of Fig. 4 shows different contexts of occurrence of the Romanian word "masa". As with any other poli-semantic words, its sense is fixed by the context. Google Translate applies statistics to disambiguate and, as can be seen, remarkably well. For the time being, I cannot imagine a workable logical solution to this issue, and if this would ever be achieved, what would be the cost of the supportive lexico-semantic resources?

But it is also clear that words induce different reactions in humans, depending of their culture, the moment of the utterance and any other pragmatic conjuncture. A notorious example is the wood language. I wonder how would a logical approach detect the humorous effect that is conveyed by phrases such as the following:

"Obiectivele majore, de însemnătate cu adevărat istorică, pe care Partidul Comunist Român le-a încredinţat frontului culturii româneşti în perioada făuririi societăţii socialiste multilateral dezvoltate angajează – deschis şi plenar…"[5] (approximately: *The major objectives,*

[5]Florian Georgescu (1982) The History Museum – dynamic factor in implementing the P.C.R. policy of patriotic education of the masses, of formation of socialist consciousness, in the Scientific Session: Permanency, unity and progress in the history of the Romanian people. Romanian Communist Party at the 60th anniversary, VI. www.mnir.ro

of truly historic significance, which the Romanian Communist Party has entrusted to the frontline of the Romanian culture in the making of the multilaterally developed socialist society – openly and fully. . .) (18)

"Pus astfel în lumină, ancorat în sinergia faptelor, recursul la universalitate nu eludează meandrele concretului."[6] (approximately: *Thus put in light, anchored in the synergy of its facts, the appeal to universality does not circumvent the meanders of the concrete.*) (19)

"Să luptăm pentru propăşirea neamului şi aducerea României pe cele mai înalte culmi de civilizaţie multilateral dezvoltată." (approximately: *Let's fight to thrive our stirps and bring Romania on the highest peaks of multilaterally developed civilization.*) (20)

3.4 Style in literature

Humans perceive co-occurrence of words as producing very suggestive images. Confronted with the extraordinary diversity of suggestion that words can convey, logic seems to me faint, forceless, impuissant. How could poetical expressions, such as the following:

"constelaţia ochilor mei" (*the constellation of my eyes*), "atingi cu auzul" (approximately: *your hearing touches*), "nisipuri de fiară" (*beast sands*) – Nichita Stănescu, *Autoportret în timp de veghe* (Auto portrait during watch time) (21)

be encoded in logical constructions? Or how could emotions incurred in sentences like the following one be seized in logical expressions?:

"It's enough for me to be sure that you and I exist at this moment." –Garcia Marquez: *One Hundred Years of Solitude* (22)

I agree that an effort to formalise in logical terms a metonymic sense of an expression, as in this magnificent sequence of simple words:

"lipindu-se de răcoarea tocului uşii" (approximately: *sticking to the chil of the door frame*) — Garcia Marquez: *One Hundred Years of Solitude* (23)

[6]attributed to Ion Iliescu

in which a touched object is replaced with a sensation that the agent borrows from that object is a challenging task. The same happens when looking for logical equivalent of metaphorical language, as here:

"Te mângâi cu degetele muiate în amintiri." (*I caress you with my fingers dipped in memories.*). (24)

4 And the solution is? (instead of conclusions)

The last examples I have given address the philosophical question of whether it is worth looking for a formalisation of language able to encode all its extremely large diversity of expressing power. Supposing logic proves to be successful in representing some language aspects, there should be a limit where the ambition to express natural language in logical form has to stop because it reaches an insurmountable limit.

Some people, including me, think that since humans use language all the time, but only rarely make explicit use of logic in their lives, the domain of NLP should not necessarily be dependent on logical formalisms. Their lack of confidence in logic as a universal machinery for processing language comes from observations of the inability of logic to support exhaustively the processing infrastructures of language. In their opinions, there are aspects of language production and understanding for which something else than logic should be used in order to explain and reproduce on the machine these human performances. On the other hand, these people agree that logic is necessary for acquiring certain types of representations. What they don't believe is that each sentence, or each sequence of sentences, should be transformed into a theorem that necessitates a proof.

It has to be clear by now that I am not talking here about reasoning, as the one needed to make volitional connections, to find explicit links, without which understanding would be impossible, as the ones exemplified in sections 2.2 and 2.3, where, clearly, logic is on the first plan, but about the primary processes that enable the use of language as a communication channel, therefore that allow cognition based on language.

However, it can be said that the whole domain of formal linguistics is inspired by logical formalisms. Indeed, why representing NL sentences as trees? Because, doing this, we intrinsically incorporate decisions about their ambiguities, this way preparing the path towards semantic representations and reasoning. A logical system applied to language means to interpret signs of the language, i.e. words, in their surface variation, as dictated by morphology, then their sequences, as dictated by syntax, and their meanings, as dictated by semantics, with the goal of arriving to an overall formal and unambiguous representation. Above the sentence boundaries, the sentential representations would be combined in discourse trees (sometimes graphs), on which rhetorical deductions could be made and inter-sentential links, as those implied by anaphorae, would be fulfilled, or extra-textual connections, as those evoked by named entities in the cultural background of the receiver of the message, would be activated. If sufficiently sensitive antennae oriented towards the external world would also be available, pragmatic contexts could make subtle revisions to these representations.

But is this enough? Suppose the day D has come when all computational theories now evolving in the field of language with the aim to decipher and represent language **in symbolic form** would reach a successful finalization, and a sufficiently rich collection of accompanying resources, necessary to support with data these formalizations, would be acquired. Are we done with the interpretation of language in that day? Can we install this tremendous computational machinery on a high performance computer or on an whatever network of cloud interconnected devices and say: from now on, whatever text we read (we, the humans), this Goliath super-computer can also read and it will get similar reactions (of course, inventoried in a very rich annotation language)?

First, let's notice that many successful NLP applications already exist, that are so dumb in a real "interpretation" of the language that would horripilate a "bad" classical linguist. Among them: machine translation. The Google Translate machine, exemplified in Fig. 4, does not "understand" a iota, in the classical sense. Put to represent the meaning of those sentences, it will be incapable. And yet, it deals so well with those sentences: the result is equivalent to that obtained by a human being graduated in Romanian-English translation.

In these approaches, of a purely statistical nature, the performance to translate any source language text into any other target language can be obtained by "compiling" a very large collection of parallel documents and a very large collection of target language documents, out of which huge tables of figures, called language models, are extracted. This computer performance in fact copies the human ability to learn a language by practising it, instead of using grammar books and drill exercises that formally incorporate the language competence.

Then, going a little bit further, we may think that a similar statistical apparatus could be used in a dialog system, that would support a human-machine dialogue, resembling intelligence (as in a Turing's test). Again, a machine, statistically trained to answer questions, could arrive to a similar level of performance as a humanly incorporated call-centre operator. Also, close to this, many models of now-a-days chat-bots make use of purely statistical solutions.

However, we should not exaggerate with congratulations and compliments. Until now, statistical solutions proved to behave well in applications where of interest is the conveying of meaning more than the stylistic expression, the clear message more than the poetical language that adorns the message with subliminal adds, in general there where language refinements that imply more than pure transmission of information, those that touch the domain of literature and art, are not involved. But, although still far away from any acceptable solutions, challenged to approach this side of language, my conviction is that statistical methods (and neural) have more chances than rule-based ones.

A vivid area of research in NLP is called Sentiment Analysis. The type of applications belonging to this domain addresses the interests that big commercial companies have to interpret opinions from their clients involving their products, in order to improve them or to estimate future trends. Big data methods put to work on text files are being employed successfully here. Still, a sentiment means much more than the mere and overtly expression of a taste or inclination. For instance, actual systems would perhaps categorise as positive a sentence like *I love you.* and as neutral one that says *I see and I smell everything around us differently since I met you.*

I am confident that machines will arrive to interpret texts, in the sense of extracting the information contained in them (if I would not believe in this success, why bothering to remain in the field?...). Moreover, more and more complex applications that put the interpretation of language at their very base will be on the market. However, I am rather reserved on the usefulness of pure logical approaches in practical NLP settings, one important reason for this being the difficulty of fuelling these systems with the amount of resources that are needed to support the complex reasoning processes. On the contrary, mixed approaches, which maculate the purity of logical approaches with statistics and/or neural models, have a much greater chance of being successful. But my optimism dilutes significantly if the border of semantic content interpretation is overpassed, and we will dig our feet on touching more subtle aspects of language, those that involve emotion sourced in language and interpretation of artistic style, therefore those that address the genuine and inspired juxtaposition of words.

References

[1] A. Cornilescu, "On the Linearization of Adjectives in Romanian," in *Romance Languages and Linguistic Theory, 2006*, D. Torck and L. Wetzles, Eds. Amsterdam: John Benjamins Publishing, 2009, pp. 45–68.

[2] A. Falaus, "Romanian N-words as Negative Quantifiers," in *Working Papers in Linguistics, Proceedings of the 31st Annual Penn Linguistics Colloquium*, (University of Pennsylvania), vol. 14, no. 1, 2008, pp. 119–134.

[3] M.-C. de Marneffe and C. D. Manning. *Stanford typed dependencies manual.* (2008). [Online]. Available: http://nlp.stanford.edu/software/dependencies_manual.pdf

[4] J. M. Moravcsik, "Aitia as Generative Factor in Aristotle's Philosophy," *Dialogue*, vol.14, no. 4, pp. 622–638, 1975. DOI: 10.1017/S001221730002655X.

[5] J. Pustejovsky, "The Generative Lexicon," *Computational Linguistics*, vol. 17, no. 4, pp. 409–441, 1991.

[6] J. Pustejovsky, *The Generative Lexicon.* Cambridge, MA: MIT Press, 1995, 312 p. ISBN: 9780262161589.

[7] J. Pustejovsky and E. Jezek. *Integrating Generative Lexicon and Lexical Semantic Resources*, LREC Tutorials, May 23 (2016). [Online]. Available: http://lrec2016.lrec-conf.org/media/filer_public/2016/05/10/tutorialmaterial_pustejovsky.pdf.

[8] S. Schoenmackers, J. Davis, O. Etzioni and D. Weld, "Learning First-Order Horn Clauses from Web Text," *Proceedings of EMNLP '10, Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, (Cambridge, Massachusetts), Stroudsburg, PA, USA: Association for Computational Linguistics, 2010, pp. 1088–1098.

[9] S. Schoenmackers, O. Etzioni and D. Weld, "Scaling Textual Inference to the Web," *Procs. of EMNLP '08, Proceedings of the Conference on Empirical Methods in Natural Language Processing*, (Honolulu, Hawaii), Stroudsburg, PA, USA: Association for Computational Linguistics, 2008, pp. 79–88.

[10] M.-D. Vulchanova, "Modification in the Balkan nominal expression: An account of the (A)NA – AN(A) order contract," in *From NP to DP. Volume 1: The syntax and semantics of noun phrases*, Martine Coene and Yves D'hulst, Eds. Amsterdam/Philadelphia: John Benjamin Publishing Company, 2003, pp. 91–118.

Dan Cristea
Received June 15, 2016

"Alexandru Ioan Cuza" University of Iaşi,
Faculty of Computer Science and
Romanian Academy, Institute for Computer Science,
E–mail: dcristea@info.uaic.ro