The Future of Text-Meaning in Computational Linguistics

Graeme Hirst

Department of Computer Science
University of Toronto
Toronto, Ontario, Canada M5s 3G4
gh@cs.toronto.edu

Abstract. Writer-based and reader-based views of text-meaning are reflected by the respective questions "What is the author trying to tell me?" and "What does this text mean to me personally?" Contemporary computational linguistics, however, generally takes neither view. But this is not adequate for the development of sophisticated applications such as intelligence gathering and question answering. I discuss different views of text-meaning from the perspective of the needs of computational text analysis and the collaborative repair of misunderstanding.

1 Introduction: Text and Text-Meaning in Computational Linguistics

In this paper, I will describe how new applications in computational linguistics and natural language processing are leading to changes in how the field views the idea of the meaning of a text. By a *text*, I mean a document or a dialogue, or a structurally complete fragment, such as a paragraph or section of a document or a turn or sequence of turns in a dialogue. Although a text could be just a single sentence, when we call it a text we add the idea that we are regarding it as complete by itself. The term *text-meaning*, then, denotes the complete in-context meaning or message of the text. It is thus potentially much more than just the sum of the sentence-meanings of the individual sentences of the text; it includes the broader "message", and possibly even a subtext as well. It will include implicatures and pragmatic and logical inferences that follow immediately from the text. A computer system (or person) might be able to understand each sentence of a text at the sentence-meaning level and yet fail to understand the text itself.

2 Where Is the Meaning of a Text?

Given this definition, we can now ask: Who decides what a text means? Where does text-meaning lie, what is its *locus*? There have been three traditional answers to this: It is the writer or speaker; it is the reader or listener; or it is the text itself.

In the first answer, that the locus of meaning is the writer or speaker, the view is that meaning is the *intent* of the agent who is communicating or, at least, is derived from that intent. A text means whatever its author intended it to mean, which might be quite different from the "literal" (sentence-level) meaning, regardless of whether the reader

P. Sojka et al. (Eds.): TSD 2008, LNAI 5246, pp. 3–11, 2008.

or listener is able to determine that intent correctly or not. This view is associated with philosophers such as Grice [8]. In the second answer, that the locus is the reader or listener, the view is that a text-meaning is each reader's individual response to, or experience of, reading the text; a reader, after all, cannot ever know a writer's intentions for certain, but can know only their own response to the text. This view, in various forms, is associated with postmodernists such as Stanley Fish and Roland Barthes, among others. In particular, Fish [6] claims that when readers agree on a meaning – that is, have the same response – it is because they are members of the same "interpretive community".

The third answer, that meaning is in the text itself, is what Reddy [23] has called the *conduit* view: that the text is a conduit through which meaning is sent from writer to reader. While Reddy argued against this as simplistic, a view of text as the locus of meaning was central to theories of literary criticism (such as that of Wimsatt and Beardsley; see Fish ([6]: 2)) that regarded the reader as too variable and the writer's intents as unknowable; only the text itself is fixed and knowable. Olson [21,22] has argued that, historically, the gradual emergence of the concept of the written *autonomous text* as a bearer of meaning was a development from the more-natural and more-fundamental concept of meaning as something that a listener adds to a spoken utterance in its context.

The question of where the meaning lies can, of course, be asked not only of texts but equally of linguistic elements at a lower level – words, sentences, semantic roles, lexical relations, and so on – with the same three potential answers. It does not follow, however, that the same answer need be chosen for each of these elements, nor that it be the same answer chosen for the text-meaning level. For example, one could argue that the effects of individual readers or of writers' intents are apparent only at the text-meaning level and not below. But equally, one could argue conversely that the idiosyncrasies of individual readers that are observed at lower levels (for example, in interpretations of lexical relations (Klebanov [14], Morris and Hirst [19], Hollingsworth and Teufel [12])) are *dampened* by constraints that arise from consideration of a text as an integrated whole, and this dampening effect is what enables texts to be autonomous bearers of meaning.

Each of the three views has its passionate defenders, but my goal here is not to argue for one over the other, but rather to regard each one as a view of meaning that is helpful in some situations in NLP.

Nonetheless, I take the view that the text itself must always be taken as a locus of meaning for all elements at all levels. (*A fortiori*, I would argue that the meaning of closed-class words is always *solely* textual, constant for all competent users of a language.) The text, after all, plays a central role in the meaning business. This is not to say that meaning cannot exist in the absence of text. Any kind of information transfer can entail meaning; one can say that a person's (non-verbal) actions have meaning and that events in the world have meaning. For example, Nadia putting the garbage bin out means that (she thinks that) it's Wednesday; or that smoke means fire (a famous example of Barwise and Perry [2]). But in the case where text is a medium of information transfer from which meaning somehow arises, one cannot overlook its causal role even if one does not accept the conduit metaphor. The question then remains as to whether the reader (listener) and/or the writer (speaker) are to be considered loci *in addition* to the text.

3 Computational Views of Text-Meaning: A Very Short History

Computational linguists are not philosophers or literary theorists, and generally don't think very much about the issues raised in the previous section. Nonetheless, any research in computational linguistics and natural language processing that involves semantics must at least implicitly choose a view of meaning. In Hirst [11], I pointed out that each of the three views of text-meaning has dominated one of the last three decades.

Ironically, in the traditional, more logic-based approaches of the mid-1970s to the mid-1980s, the dominant view was the most postmodern one, the reader-based view: language understanding was viewed as a knowledge-based enterprise, and so a text was understood by a system in the light of its own particular knowledge. This was so at both the sentence-meaning level and the text-meaning level; in the latter case, exemplified by the work of Schank and his colleagues (e.g., Schank [24], Schank and Abelson [25]), to understand language was to use knowledge and inference to "fill in the gaps" in what was said. The reader-based view became most explicit in Corriveau's computational model of time-constrained text comprehension (published [4], but developed in the late 1980s).

In the subsequent decade, with a greater interest in interactive dialogues, attention became focused on the speaker and their intentions. Again, this was so at both levels. At the text-meaning level, the work was typified by Carberry's book *Plan Recognition in Natural Language Dialogue* [3]. And below the text-meaning level, Grosz and Sidner [9], for example, took a speaker- or writer-based view of discourse structure.[1]

From the mid-1990s into the 2000s, with the increasing availability and importance of autonomous text such as newswire and Web pages and the rise of statistical and machine-learning methods, attention became focused solely on the text as the locus of meaning, with no thought or individuation of the writer or the reader. Tasks that implicitly entailed a regard for text-meaning, such as summarization, machine translation, and topic tracking, became seen just as meaning-preserving statistical transformations of the text. Text was to be "processed", not understood. In fact, in this paradigm, text-meaning was taken to be little or nothing more than the sum of sentence-meanings; for example, a text could be summarized, in this view, just by finding and displaying its "most important" sentences.

4 Computational Views of Text-Meaning: The Future

With this as background, I want to now turn to future views of text-meaning in computational linguistics and the future role of *the linguistic computer*. We are now starting to again see applications in natural language processing in which writer- and reader-based views of text-meaning are more explicit in the task, and I believe that, as NLP methods advance, these kinds of application will become increasingly important. I will give examples below.

In fact, the word *processing* in *natural language processing* will become less appropriate, as these developments will bring us closer to what early researchers optimistically

[1] Contrary to remarks by Morris [18], the listener or reader played no role in determining the discourse structure or focus of attention in Grosz and Sidner's theory; see Section 4.3 below.

called *natural language understanding*. But *understanding* is not the right word either; it's too loaded, and tends to imply a single "correct" meaning. A better word would be *interpretation*, implying one of many possible kinds of understandings. And it's exactly in the distinctions between different kinds of understanding or interpretations that these applications have their value.

In the discussion below, I will divide applications into two classes: Those in which the computer is an **observer** of external text, playing only the role of reader; and those involving interactive dialogue in which the computer is an **active participant** and hence plays alternating roles of reader and writer or speaker and listener. In the first case, the text (which could be a dialogue between others) has been created for some other reader in some other context; typically, the computer will be looking at it on behalf of a user in order to achieve some goal for that user. In the second case, the text that the computer reads is created by the writer expressly for the computer in the context of the dialogue; typically, this is a user interacting with a computer in natural language in order to achieve some goal.[2] In the space available here, I'll concentrate mostly on the first case, observing external text.

4.1 What Does This Text Mean to Me?

One of the profound changes to everyday life that has resulted from the rise of the search engine and the World Wide Web over the last decade is that the average person has become a researcher. In addition to students using Web research in their school assignments, people are now turning to Google and the Web to answer questions arising in their daily life that they previously would not have bothered about, and usually their starting point is a search engine.[3] Questions include those about health matters, finding a place to live, hobbies, government services, and matters of general curiosity. But of course, such questions must be re-expressed by the user as search terms, and the resulting documents must be read to extract the answer – assuming that they do, in fact, contain the answer.

Now, observe that a document might happen to provide an answer to a user's particular question without the writer having had any intent of doing so. But when search is based on keyword matching, this will be relatively infrequent; in the canonical case, a relevant document contains the search terms from the question[4] (and probably in close proximity if it is highly ranked by Google) exactly because the writer has made an explicit statement that answers the question. However, the more abstract, wide-ranging, or unusual the question, the less likely this is, and indeed the more likely it is that a relevant document will not be a good keyword match at all. These characteristics can be seen, for example, in the test topics used in the "HARD" task of the Text Retrieval Conference (TREC; `trec.nist.gov`) (Allan [1]); for example, "What factors contributed to the

[2] Of course, the two classes might occur together in a single application; the user might have a dialogue with the computer in order to achieve a goal pertaining to some observed texts.

[3] This is the general trend in the United States shown by recent surveys by the Pew Internet & American Life Project (`http://www.pewinternet.org`). There is no reason to think that things are significantly different in other developed countries.

[4] Google's PageRank algorithm may return documents that do not contain the search terms, but this is a minor effect for this argument.

growth of consumer on-line shopping?" In TREC, the requirements of this task were merely document or passage retrieval; but the task more generally can be thought of as the evolution and confluence of several streams of research in NLP and information retrieval – document retrieval, question-answering, and multi-document summarization – and future systems (we anticipate) will construct a full and relevant answer to the user's question by selecting relevant information from many documents. To do so, such a system must consider each document or passage from the point of view of the user's question (and anything else known about the user);[5] that is, it must take a reader-based view of the meaning of the text: "What does this text mean to me?"

Another particular research direction of relevance here is that known as "learning by reading" (LbR) (Hovy [13]). Much as in the research of the 1970s–80s mentioned in Section 3 above, the goal is for the system to achieve a full understanding of a (possibly lengthy) text, beginning from an initial knowledge base (which might be quite large), making inferences and interpretations, and then answering questions about what it has read. And like that earlier work, because of its dependence on prior knowledge for the interpretation of new text, it takes a reader-based view of text-meaning (but see also Section 4.2 below); and if its prior knowledge is deficient, it may misinterpret text (Forbus et al. [7]). Moreover, in the context of LbR, Ureel *et al* [28] and Forbus et al. [7] present a prototype of a "ruminator" process that "reflects" on what it has read by using its prior knowledge to generate questions that, if answered, would fill gaps in the system's understanding of the text.[6]

4.2 What Are They Trying to Tell Me?

The complement of applications that take a reader-based view of text is applications that try to determine the intent of the author of the text, and, to the extent that it's possible, without the bias or influence of any reader-based information. That is, the goal is to figure out "what is the writer trying to say?" Such applications can be broadly characterized as a kind of *intelligence gathering* – not necessarily in the military or espionage sense, though those domains are canonical examples, but in the broader sense of trying to find out, for any reason, what other people believe, or are thinking, or are planning.

A good example is the development in the last few years of computational methods for *sentiment analysis*, especially for blogs and similar online texts, and the more general idea of *opinion analysis, extraction,* or *mining* that has emerged in the NLP literature. The goal is to discover, from evaluative opinions expressed online, what consumers really think about various commercial products and services; companies can use this information in planning their market strategies. Clearly, if such information is to be maximally useful, it must be a close reflection of the writer's opinion, uncolored by any individualities or preconceptions of the reader. Work on tasks such as multi-perspective question answering (Stoyanov, Cardie, and Wiebe [26]) has extended this research into

[5] This could be facilitated by a personalized lexicon or thesaurus based on the user's past reading, as proposed by Yoshida et al. [29].

[6] Ureel et al. describe rumination as an "offline" process; hence the reflections or questions themselves can't be considered part of the text-meaning. But by its nature the process is nonetheless closely related to a reader-based view of text-meaning.

finding and summarizing reported or expressed opinions in newswire text. But beyond this, we can expect sophisticated applications of the future to be able to find implicitly stated opinion and sentiment without restriction to any pre-defined topic. Eventually, this may use (possibly inferred) knowledge of the writer and the writer's prior beliefs to determine just what the writer intends to say; that is, systems may be able to classify texts by the ideological background of the writer and interpret and analyze them accordingly. (This goes well beyond Malrieu's [17] description of a computational method of checking texts for *ideological consistency*.)[7]

Learning by reading will eventually come under this heading too, because as defined by Hovy [13], its goal of being able to answer questions about the text is construed as *test-taking*: in the prototypes, the text is taken from a high-school book on chemistry and the questions are from an advanced-placement exam on the material. Thus there is a "right answer" to the questions, and despite the reader-based view of understanding that the LbR architectures entail (see Section 4.1 above), LbR nonetheless has the writer-based goal of figuring out "what are they trying to tell me" in order to learn enough to pass the exam. However, this is for the future; there is no recognition of the writer or of their intent in the LbR architectures presently proposed (Hovy [13], Forbus et al. [7]).

Last, it may be recalled that the goal of high-quality machine translation is, by definition, determining, and preserving across languages, an author's intent – a fact that is largely obscured by the still-dominant purely statistical text-transformation approach. As the renewed interest in interlingual, and hence semantic, methods in MT develops further (e.g., Nirenburg and Raskin [20], Farwell et al. [5]), MT too will take on a more-explicit writer-based view of text-meaning.

4.3 Recovering from Misunderstanding

Despite any latitude a reader has in interpretation in all that we have said above, an interpretation can nonetheless count as a *misunderstanding* or *misinterpretation* if it does not fully respect the given text. That could result from a simple mishearing or misreading of the text, or from inattention or mental confusion; but it could also result from linguistic processing that is at odds with what the writer or speaker expected of the reader – incorrect resolution of an ambiguity. Examples of this include resolving an anaphor to a different antecedent than intended, choosing a different sense of a homonym, and attaching a phrase to a different point in the parse tree.

Although there is no guarantee that a misunderstanding will be detected as such, a listener or reader might hypothesize a misunderstanding, either at the present point in the text or at some earlier point, if he or she cannot interpret the present text. A speaker might hypothesize a misunderstanding by the other if they respond in an unexpected way. In either case, recognition that a misunderstanding has occurred might be followed

[7] Some of what I earlier suggested were purely text – view-based applications might also be considered to be implicitly about opinion analysis. For example, in processing newswire text about current objective events, while there might be said to be no overt authorial intent other than to comprehensively communicate facts, often, even in "objective" news reporting, an authorial bias or intent or ideology may nonetheless be implied in what facts are communicated and what words are used to describe them; such an intent might be discovered by comparing two authors' reports of the same event.

by recognition of the likely misunderstanding itself, leading to a reinterpretation by the reader or listener in the first scenario, or to a clarification by the speaker in the second scenario.

Given that computers will remain relatively poor understanders for some time yet, it is important that they be able to detect misunderstanding in dialogue, both their own and that of their conversant. Some years ago, McRoy ([15], McRoy and Hirst [16]) developed computational models of negotiation in the collaborative, constructive repair of misunderstanding in dialogues. Reflecting the inherent symmetry of the negotiation of meaning, the model could play both the role of the conversant who is misunderstood and the role of the conversant who fails to understand.

For example, the model could account for the text-meaning-level misunderstanding in this fragment of a conversation between a mother and her child Russ about a forthcoming parent-teacher meeting (Terasaki [27]):

1. MOTHER: Do you know who's going to that meeting?
2. RUSS: Who?
3. MOTHER: I don't know.
4. RUSS: Oh. Probably Mrs McOwen and some of the teachers.

Russ initially interprets line 1 as expressing Mother's desire to tell, that is, as a *pretelling* or *pre-announcement* as if Mother intends to surprise him (cf *Guess who's going to that meeting!*). But Russ finds this interpretation inconsistent with her next utterance; in line 3, instead of telling him who's going, as he would expect after a pretelling, Mother claims that she does not know. Russ recovers by reinterpreting line 1 as an indirect request for information, which his line 4 then responds to. In this example, we see that both the speaker and listener negotiate and refine the meaning of a prior utterance. (They swap roles of speaker and listener in the dialogue as they do, but not their roles with respect to the original utterance.) McRoy's model is thus a rare example in computational linguistics of one in which both speaker and hearer are loci of meaning. It thus contrasts with, for example, the earlier work of Grosz and Sidner [9] on discourse segmentation in which, even in the case of interactive dialogue, nothing is negotiated or refined and only the speaker, not the listener, is a locus of meaning; each participant may modify the structure or focus of the discourse only when taking their turn to act as speaker; and when the other takes their turn as speaker, they cannot change or undo what the first has done.

5 Conclusion

There are many other aspects of how we view meaning in computational linguistics and natural language processing and how that will change in the future that I haven't had space to touch on in this paper. These include the negotiation or collaborative construction of meaning in interactive dialogue, particularly in the interactive elicitation of knowledge; methodological issues that arise within annotation-based learning paradigms; the idea of current research in textual entailment as a preliminary to the more-general task of searching for differing interpretations of a text. This in turn could lead to automatic processes for reconciling seemingly incompatible interpretations of the same text or situation: the linguistic computer as an aid to mediation and reconciliation.

Acknowledgments. Some of the ideas in this paper are elaborations or improvements on points made in Hirst [10] and Hirst [11]. I am grateful to Stephen Regoczei and Jean-Pierre Corriveau for first introducing me to the ideas of reader-based models of meaning; to Susan McRoy, Peter Heeman, Phil Edmonds, and David Traum for discussions of models of misunderstanding; and to Jane Morris for drawing my attention to how the view in CL of the locus of text-meaning differed in the early 2000s from that of the 1980s. I thank Nadia Talent for helpful comments on an earlier draft of this paper. The preparation of this paper was supported by a grant from the Natural Sciences and Engineering Research Council of Canada.

References

1. Allan, J.: HARD track overview in TREC 2005. In: The 14th Text REtrieval Conference (TREC 2005) Proceedings, NIST (2006)
2. Barwise, J., Perry, J.: Situations and Attitudes. MIT Press, Cambridge (1983)
3. Carberry, S.: Plan Recognition in Natural Language Dialogue. MIT Press, Cambridge (1990)
4. Corriveau, J.-P.: Time-constrained Memory: A reader-based approach to text comprehension. Lawrence Erlbaum Associates, Mahwah (1995)
5. Farwell, D., et al.: Interlingual annotation of multilingual text corpora and FrameNet. In: Boas, H. (ed.) Multilingual FrameNets in Computational Lexicography, Mouton de Gruyter (to appear)
6. Fish, S.: Is there a text in this class? The authority of interpretive communities. Harvard University Press (1980)
7. Forbus, K.D., et al.: Integrating natural language, knowledge representation and reasoning, and analogical processing to learn by reading. In: Proceedings, 22nd AAAI Conference on Artificial Intelligence (AAAI-2007), Vancouver, pp. 1542–1547 (2007)
8. Grice, H.P.: Utterer's meaning, sentence-meaning, and word-meaning. Foundations of Language 4, 225–242 (1968)
9. Grosz, B.J., Sidner, C.L.: Attention, intentions, and the structure of discourse. Computational Linguistics 12(3), 175–204 (1986)
10. Hirst, G.: Negotiation, compromise, and collaboration in interpersonal and human-computer conversations. In: Proceedings, Workshop on Meaning Negotiation, 18th National Conference on Artificial Intelligence (AAAI-2002), Edmonton, pp. 1–4 (2002)
11. Hirst, G.: Views of text-meaning in computational linguistics: Past, present, and future. In: Dodig-Crnkovic, G., Stuart, S. (eds.) Computation, Information, Cognition – The Nexus and the Liminal, pp. 270–279. Cambridge Scholars Publishing (2007)
12. Hollingsworth, B., Teufel, S.: Human annotation of lexical chains: Coverage and agreement measures. In: Workshop on Methodologies and Evaluation of Lexical Cohesion Techniques in Real-world Applications, Salvador, Brazil (2005)
13. Hovy, E.: Learning by reading: An experiment in text analysis. In: Sojka, P., Kopeček, I., Pala, K. (eds.) Text, Speech and Dialogue. LNCS (LNAI), vol. 4188, pp. 3–12. Springer, Heidelberg (2006)
14. Klebanov, B.B.: Using readers to identify lexical cohesive structures in texts. In: Proceedings, Student Research Workshop, 43 rd Annual Meeting of the Association for Computational Linguistics, Ann Arbor, pp. 55–60 (2005)
15. McRoy, S.: Abductive interpretation and reinterpretation of natural language utterances. Ph.D. thesis, Department of Computer Science, University of Toronto (1993)
16. McRoy, S., Hirst, G.: The repair of speech act misunderstandings by abductive inference. Computational Linguistics 21(4), 435–478 (1995)

17. Malrieu, J.P.: Evaluative Semantics. Routledge (1999)
18. Morris, J.: Readers perceptions of lexical cohesion in text. In: Proceedings of the 32nd annual conference of the Canadian Association for Information Science, Winnipeg (2004)
19. Morris, J., Hirst, G.: The subjectivity of lexical cohesion in text. In: Shanahan, J.G., Qu, Y., Wiebe, J. (eds.) Computing attitude and affect in text, pp. 41–48. Springer, Heidelberg (2005)
20. Nirenburg, S., Raskin, V.: Ontological Semantics. MIT Press, Cambridge (2004)
21. Olson, D.R.: From utterance to text: The bias of language in speech and writing. Harvard Educational Review 47(3), 257–281 (1977)
22. Olson, D.R.: The World on Paper. Cambridge University Press, Cambridge (1994)
23. Reddy, M.J.: The conduit metaphor: A case of frame conflict in our language about language. In: Ortony, A. (ed.) Metaphor and Thought, pp. 284–324. Oxford University Press, Oxford (1979)
24. Schank, R.C. (ed.): Conceptual Information Processing. North-Holland, Amsterdam (1975)
25. Schank, R.C., Abelson, R.P.: Scripts, Plans, Goals and Understanding. Lawrence Erlbaum Associates, Mahwah (1977)
26. Stoyanov, V., Cardie, C., Wiebe, J.: Multi-perspective question answering using the OpQA corpus. In: Proceedings, Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP), Vancouver, pp. 923–930 (2005)
27. Terasaki, A.: Pre-announcement sequences in conversation. Social Science Working Paper 99. University of California, Irvine (1976)
28. Ureel II, L., et al.: Question generation for learning by reading. In: Proceedings of the AAAI Workshop on Inference for Textual Question Answering, Pittsburgh, pp. 22–26 (2005)
29. Yoshida, S., et al.: Constructing and examining personalized cooccurrence-based thesauri on Web pages. In: Proceedings, 12th International World Wide Web Conference, Budapest (2003)