

Introduction to the Special Issue on Language in Social Media: Exploiting Discourse and Other Contextual Information

Farah Benamara
Paul Sabatier University
IRIT-Université de Toulouse
benamara@irit.fr

Diana Inkpen
University of Ottawa
School of Electrical Engineering and
Computer Science
Diana.Inkpen@uottawa.ca

Maite Taboada
Simon Fraser University
Department of Linguistics
mtaboada@sfu.ca

Social media content is changing the way people interact with each other and share information, personal messages, and opinions about situations, objects, and past experiences. Most social media texts are short online conversational posts or comments that do not contain enough information for natural language processing (NLP) tools, as they are often accompanied by non-linguistic contextual information, including meta-data (e.g., the user's profile, the social network of the user, and their interactions with other users). Exploiting such different types of context and their interactions makes the automatic processing of social media texts a challenging research task. Indeed, simply applying traditional text mining tools is clearly sub-optimal, as, typically, these tools take into account neither the interactive dimension nor the particular nature of this data, which shares properties with both spoken and written language. This special issue contributes to a deeper understanding of the role of these interactions to process social media data from a new perspective in discourse interpretation. This introduction first provides the necessary background to understand what context is from both the linguistic and computational linguistic perspectives, then presents the most recent context-based approaches to NLP for social media. We conclude with an overview of the papers accepted in this special issue, highlighting what we believe are the future directions in processing social media texts.

Submission received: 10 September 2018; accepted for publication: 10 September 2018.

doi:10.1162/coli.a.00333

1. Introduction

Social media content has, for many people and organizations, changed the way we interact and share information. This content (ranging from blogs, fora, reviews, and various social networking sites) has specific characteristics that are often referred to as the five V's: volume, variety, velocity, veracity, and value.

Social media texts are more difficult to process than traditional texts because of the nature of the social conversations—posted in real-time. The texts are unstructured and are presented in many formats and written by different people in many languages and styles. Typographic errors are common, and chat and in-group slang have become increasingly prevalent on social networking sites like Facebook and Twitter.

In addition, most social media texts are short online conversational posts or comments that do not contain enough information for natural language processing (NLP) tools. They are often accompanied by non-linguistic contextual information, including meta-data such as the social network of each user and their interactions with other users. Because the conversation flow is not necessarily sequential, as users can write (and hence reply) at different times, these conversations are often called **asynchronous**.

Exploiting this kind of contextual information and meta-data could compensate for the lack of information from the texts themselves. Such rich contextual information makes the automatic processing of social media content a challenging research task. Indeed, simply applying traditional text mining tools is clearly sub-optimal, as it takes into account neither the interactive dimension nor the particular nature of these data, which share properties with both spoken and written language. Most research on NLP for social media focuses primarily on content-based processing of the linguistic information, using lexical semantics (e.g., discovering new word senses or multi-word expressions) or semantic analysis (opinion extraction, irony detection, event and topic detection, geo-location detection) (Aiello et al. 2013; Ghosh et al. 2015; Inkpen et al. 2015; Londhe, Srihari, and Gopalakrishnan 2016).¹ Other research explores the interactions between content and extra-linguistic or extra-textual features, showing that combining linguistic data with network and/or user context improves performance over a baseline that uses only textual information. For example, user profiles like age, gender, and location can be used to enhance subjectivity detection (including sentiment and emotion) (Volkova, Coppersmith, and Van Durme 2014; Volkova and Bachrach 2016), vote predictions (Persing and Ng 2014), or language identification (Saloot et al. 2016). Also, information from the conversational thread structure (e.g., links between previous posts) or valuable external sources can serve as contextual constraints to better capture the sentiment or the figurative reading of an utterance (Mukherjee and Bhattacharyya 2012; Karoui et al. 2015; Wallace, Choe, and Charniak 2015)². Finally, the social network, like social relationships, can enable grouping users according to specific communities regarding the topics or the sentiments they share (Deitrick and Hu 2013; West et al. 2014).

Besides social media processing, the interaction of contextual information derived from sentences, discourse, and other forms of linguistic and extra-linguistic information have shown their effectiveness in language technology in general (Taboada and Mann 2006; Webber, Egg, and Kordoni 2012). This shows that computational linguistics is

1 See Farzindar and Inkpen (2017) for an overview of the main NLP approaches for social media.

2 See Benamara, Taboada, and Mathieu (2017) for a recent overview of context-based approaches to evaluative language processing.

currently experiencing a *discourse turn*, a growing awareness of how multiple sources of information, and especially information from context and discourse, can have a positive impact on a range of computational applications. This turn is particularly notable in the research community, where several workshops have been recently organized in major NLP international conferences to account for the role discourse and context can have in various NLP tasks (e.g., the DiscoMT series on discourse in machine translation, CompPrag on computational pragmatics, SocialNLP on NLP for Social Media, and many of the papers at *SEM or SemEval workshops).

This special issue invited contributions that implement such approaches, but not restricted exclusively to applications in evaluative language and sentiment analysis.

Before giving an overview of the papers accepted in this special issue (Section 4), we provide some background on what *context* is from both the linguistic and computational linguistic perspectives (Section 2). We then focus on current context-based approaches to NLP for social media (Section 3). We end this introduction by highlighting what we believe are the future directions in processing social media texts.

2. Context in Computational Linguistics

Context is a pervasive term in linguistics and no single coherent definition of context is available (Bach 1997; Recanati 2008; Jaszczolt 2012; Korta and Perry 2015). An intuitive view is to consider the distinctions between the *linguistic information* formed by morphological, syntactic, or textual material surrounding a word, and any other *contextual information* surrounding the utterance. Bunt and Black (2000) discuss the following non-exhaustive aspects of contextual information:

- Discourse context: What has been said before in the conversation (i.e., objects that have been introduced in the preceding discourse).
- Attitudinal or epistemic context: This encompasses the speaker's knowledge, the hearer's knowledge, and the common ground (i.e., what is known to both the speaker and the hearer about the domain of the discourse).
- Spatio-temporal properties of the situation in which the utterance occurs, like the relative time and place of speaking.
- Physical and perceptual context: Objects that are known to be present or visible in the speaker's and the hearer's environment; actions and events perceivable in that environment. The textual form of an utterance (such as punctuation and layout) is also important.
- Social context: The social relationship of the people involved in communication. A sentence like *President, leave me alone* is only shocking because we know one does not usually address a president this way.

The question is then: How can these different sources of information interact to make computers understand natural language texts? There are two possible options to answer this question: Consider each source of information as a separate stage, involving a linear process starting with words and ending with extra-linguistic context; or incorporate contextual information at an earlier stage. The first option being computationally inefficient due in particular to the ambiguity of words and sentences when processed in

isolation, this special issue adopted the second option, as explained in the subsequent sections.

2.1 Words and Sentences

One way to compute the meaning of a text is to exploit the meanings of words and how these words are syntactically composed to form a text. This inspired the development of truth-conditional semantics or model-theoretic semantics in which the meaning of a sentence is determined relative to a model, which can be taken to be an abstract description of the world (Montague 1974; Tarski 1983). Lexical meaning and syntax provide *linguistic knowledge* and play a crucial role in studying the behavior of semantic phenomena bound at the sentence level (Bos 2011).

We illustrate the composition process by the effect intensifiers and downtoners have on the evaluative expressions they modify. Many devices intensify by changing the intensity of an evaluative word, whether by bringing it up or down. For instance, adjectives may intensify or downtone the noun they accompany (e.g., *A definite success*), as adverbs do with adjectives (e.g., *A very dangerous trip*) or verbs (e.g., *He behaved badly*). Examples (1) and (2), extracted from the CASOAR corpus (Benamara et al. 2016), show a more complex case where the overall sentiment orientation is determined in a bottom-up fashion.

- (1) The actors are not good enough.
- (2) This restaurant proposes good quality Greek cuisine in a warm atmosphere.

Moving from a subjectivity lexicon that encodes the meaning of sentiment-relevant words (like the adjectives *good* and *warm*), composition follows the syntactic tree up to the main clause by combining pairs of sister nodes by means of a set of sentiment composition rules. In Example (1), sentiment calculation has first to deal with the composition *good enough* that softens the positivity of the evaluation, which in turn has to be composed with the negation (*not*) that makes the overall opinion negative. In Example (2), the sentence's syntactic structure indicates that the atmosphere and the cuisine have both a positive evaluation. For more discussions on sentiment composition, the reader can refer to the Stanford Sentiment Treebank (Socher et al. 2013).

The composition process assumes that the interpretation of a given word within a sentence is fixed or disambiguated before being combined, which makes it restrictive in that it "precludes nonlinguistic information to go into the computation of meaning" (Bunt 2001).³ Indeed, the meaning of a sentence is closely tied to the pragmatics of how language is used, and thus to the meaning of the words themselves, which can be assigned different possible readings in different situations (Pustejovsky 1995; Lenci 2006). Consider the problem of lexical ambiguity. For example, *A sad movie* expresses a sentiment or feeling of grief, whereas *Sad weather* expresses an undesirable judgment that can be paraphrased as *The weather is bad*. There are also ambiguities that are not caused by lexical choice, but by the context in which the words occur. For instance, the adjective *long* may denote a negative sentiment in restaurant reviews (cf. Example (3)) but a positive sentiment in phone reviews (cf. Example (4)). The same adjective can also be purely factual, as in Example (5).

³ See Janssen (2001) and Zimmermann (2013) for a discussion of the principle of compositionality.

- (3) There is a long wait between courses.
- (4) The smart phone has a long battery life.
- (5) It has rained for a long time.

The assumption that word meaning is a function of the contexts in which it occurs within the sentence is at the center of the distributional semantics hypothesis (Turney and Pantel 2010). Distributional models represent words by vectors build by extracting co-occurrences statistics from large corpora, then use linear algebra as a computational tool to project lexical vectors to phrase vectors. Vectorial representations are extremely effective for computing semantic similarity between words, and more generally investigating the interplay between meaning and contexts (Lenci 2018).

The meaning of a sentence can also rely on other types of information, such as prosodic information in the case of spoken utterances; or punctuation, layout, and emojis in the case of textual utterances. The latter is of particular importance when analyzing social media, as shown in Examples (6) and (7), where capitalization and character repetition, respectively, emphasize the positive opinion towards the movie.

- (6) This movie was AMAZING.
- (7) This movie was amaaazzzzzing.

2.2 Beyond Sentences: Discourse Structure

Words and sentences do not occur in isolation, but both are always part of a coherent and cohesive structure in which the discourse units are related to each other. Coherence refers to the logical structure of the discourse, where every part of a text has a function, a role to play, with respect to other parts in the text (Taboada and Mann 2006). Coherence has to do with semantic or pragmatic relations among units to produce the overall meaning of a discourse (Hobbs 1979; Mann and Thompson 1988; Grosz, Joshi, and Weinstein 1995). The impression of coherence in text (that it is organized, that it hangs together) is also aided by cohesion, the linking of entities in discourse (Halliday and Hasan 1976). Linking across entities happens through grammatical and lexical connections such as anaphoric expressions and lexical relations (synonymy, meronymy, hyponymy) appearing across sentences.

Theories of discourse interpretation typically account for meaning *beyond the sentence*. Roughly, two main approaches have been developed: dynamic semantics (Heim 1982; Kamp and Reyle 1993) and theories of discourse structure (Hobbs 1979; Grosz and Sidner 1986; Mann and Thompson 1988; Asher and Lascarides 2003; Prasad, Webber, and Joshi 2014).

The first approach extends model-theoretic semantics to account for the semantic contribution that a sentence makes to a discourse in terms of a relation between an input context prior to the sentence and an output one. Discourse context is therefore a dynamic concept:

When a sentence S is interpreted within the discourse context K , the result of its interpretation will be integrated into K . The updated context K' , which reflects the contribution made by S as well as those made by the sentences preceding it, will then be the discourse context for the next sentence. (Kamp and Reyle, 2010, page 3)

In the second approach, theories of discourse structure derive meaning from the rhetorical relations that link discourse units⁴ such as ELABORATION, EXPLANATION, NARRATION, and so forth. Discourse relations are important factors that make a discourse coherent. Coherence can be accounted for by positing relations between clauses, sentences, or speech acts (see the next section) that organize the writer's intentions (with explanations, elaborations, and contrasts, for instance) or explain speakers' turns (e.g., answer to a question, acknowledgment of a proposal or an assertion, correction of an assertion). A number of theories of relational coherence have been proposed, for written text and dialogue, which make different assumptions about the kinds of relations (thus yielding different taxonomies of discourse relations), or the resulting structure (a chain, a tree, or diversely constrained types of graphs that influence the interpretation process) (see Asher and Lascardes 2003; Taboada and Mann 2006 for an overview).

Even if dynamic semantics and theories of discourse structure differ in their aims and methods, they stress the need to model the cumulative nature of discourse interpretation, namely, the interpretation of a current discourse unit depends on the content of the part of the discourse which precedes it. To illustrate the importance of discourse structure and how constraints on coherent discourse determine lexical sense disambiguation, consider the following two short texts, taken respectively, from TripAdvisor and Twitter.⁵

- (8) [This restaurant is not remarkable.] π_1 [The dishes were correct] π_2 [but side dishes very average.] π_3 [The wine was warm.] π_4
- (9) I want to be an ecologist, but energy-saving light bulbs take more time to burst these idiots moths.

Example (8) shows that sentiment is a semantic scope phenomenon governed by discourse structure (Polanyi and van den Berg 2011). In the first sentence, the author introduces the main topic of the discourse (*This restaurant*), expressing a negative opinion towards it. This opinion is further elaborated in the discourse units π_2 to π_4 , where the author comments on two aspects of the restaurant: the cuisine and wine. To infer the ELABORATION relation that holds between π_1 and (π_2 - π_3) and between π_1 and π_4 , we need detailed lexical knowledge and probably domain knowledge as well (the fact that cuisine and wine are part of a restaurant is implicit). π_4 expresses a negative opinion lexicalized by the adjective *warm*. The interpretation of the degree of subjectivity of this adjective is a matter of context. The fact that π_4 elaborates on π_1 helps disambiguating the sense of this adjective: one cannot elaborate positively on a topic that has been previously assigned a negative opinion.

Finally, Example (9) shows the importance of discursive contextual phenomena at the sentence level: It is the contrast rhetorical relation triggered by the discourse connective *but* that allows us to infer that the writer implicitly says that they are against saving energy, even though they state the contrary in the first sentence.

2.3 Beyond What Is Said

Full comprehension of a text also requires understanding more than what is linguistically encoded, that is, understanding *beyond what is said*. Approaches like speech act

4 Some theories do also provide a model-theoretic semantics for a discourse. For instance, the Structured Discourse Representation Theory (Asher and Lascardes 2003) incorporates, but also extends, dynamic semantics.

5 This is a French tweet translated to English.

theory (Austin 1962; Searle 1969) and conversational implicature (Grice 1975) make a clear distinction between what is said by an utterance and what is *implicated* or *performed* in a particular linguistic and social context or by saying something (Korta and Perry 2015).

Austin (1962) provided a framework for connecting the literal meaning of an utterance with its intended meaning. He argued that every utterance has three layers of meaning: (i) a locutionary act that corresponds to the act of saying something with words, (ii) an illocutionary act, which conveys the speaker's intended meaning on the basis of the existence of a social practice, conventions, or "constitutive" rules in doing things with words (like ordering, offering, warning, promising, etc.), and (iii) a perlocutionary act that reflects the listener's perception of the speaker's intended meaning, that is, the effect a locutionary act has on the feelings, thoughts, or actions of either the speaker or the listener (like inspiring, amusing, persuading, etc.). For example, the illocutionary act of the utterance *I am free next week, shall we meet on Friday?* is a suggestion, while its intended perlocutionary effect might be to invite the hearer to fix a particular day to meet. The illocutionary act is a central aspect of the speech-act theory, developed later by Searle (1969).

Speech acts are the semantic/pragmatic counterpart of sentence types. The sentences types affirmative, interrogative, and exclamative correlate with the speech acts of assertion, question, expression, and order. Speech acts are relevant in social media and there is an emerging new interest in the computational community for speech acts (see, e.g., the article by Joty and Mohiuddin in this special issue).

Whereas speech acts have traditionally been understood as unary properties of expressions that convey propositions, Searle lists categories of speech acts like "answers" that are clearly relational (an answer is an answer to a particular question). Once one observes that some speech acts are relational, it is relatively straightforward to see discourse relations like EXPLANATION and ELABORATION also as types of speech acts. Unlike traditional speech acts, however, instances of discourse relations easily embed under various operators (like modality), whereas it remains controversial as to whether speech acts like assertion or requests embed.⁶

Speech acts are crucial in the analysis of some pragmatic phenomena such as preferences and intentions that concern the future states of affairs or plans that one wants to achieve. For example, in the conversational thread for Example (10) (taken from Twitter), the question-answer pair that links User's A question to User's B answer helps to better capture User B's intention towards eating organic food and not food with additives or pesticides.

(10) (User A) Do you prefer eating cakes with additives or fruits with pesticides?

(User B) Neither. I prefer to eat organic.

On the other hand, Grice (1975) argued that communication between people was also characterized by the process of intention recognition. He made a clear distinction between what is said by an utterance (i.e., meaning out of context) and what is implied or meant by an utterance (i.e., meaning in context). In his theory of conversational implicature, Grice proposes that to capture the speaker's meaning, the hearer needs to rely on the meaning of the sentence uttered, contextual assumptions, and the Cooperative Principle, which speakers are expected to observe. The Cooperative Principle states that speakers make contributions to the conversation that are cooperative, and is

⁶ See the work of Krifka (2002) for arguments that even standard speech acts embed to some degree.

expressed in four maxims that the communication participants are supposed to follow. The maxims ask the speaker to say what they believe to be the truth (Quality), to be as informative as possible (Quantity), to say the utterance at the appropriate point in the interaction (Relevance), and in the appropriate manner (Manner). The maxims are, in a sense, ideals, and Grice provided examples of violations of these maxims for various reasons. The violation of a maxim may result in the speaker conveying, in addition to the literal meaning of the utterance, a meaning that does not contribute to the truth-conditional content of the utterance, which leads to conversational implicature. Implicatures are thus inferences that can defeat literal and compositional meaning. Example (11) is a typical example of relevance violation: B conveys to A that they will not be accepting A's invitation for dinner, although they have not said so directly.

- (11) A. Let's have dinner tonight.
 B. I have to finish my homework.

Grice makes the important assumptions that participants in a discourse are rational agents and that they are governed by cooperative principles. However, in some cases involving non-literal readings or negotiation, agents do not always have rational communicative behavior.

Some contemporary researchers reject the distinction between literal and utterance meaning, arguing that what is said is always dependent on the context (Recanati 2004; Korta and Perry 2015). The debate shared by literalists and contextualists on the frontier between semantics and pragmatics is not the most important point here.⁷ What matters for the purpose of this special issue is how to make computers capture the meaning of a text when immersed *in the context in which it is uttered*.

In user-generated content such as product reviews, inference is often needed to capture implicit evaluation like the ones expressed in the movie reviews of Examples (12) and (13), taken from the CASOAR corpus. Even if there are no explicit subjective words, everyone would expect a movie to be good when reading Example (12), and bad after reading Example (13).

- (12) This is a definite choice to be in my DVD collection.
 (13) I really want my money back.

Irony is another important pragmatic phenomenon that poses new challenges when processing short texts. Irony can be defined as an incongruity between the literal meaning of an utterance and its intended meaning (Grice 1975; Sperber and Wilson 1981; Utsumi 1996; Attardo 2000). In social media, such as Twitter, and mainly in English, users apply specific hashtags (#irony, #sarcasm, #sarcastic) to help readers understand that a message is ironic. This is shown in the tweet of Example (14), which clearly expresses a negative opinion towards Nabilla, although there are two positive opinion words (*classy* and *beautiful*).

- (14) #Nabilla a very classy and beautiful girl, not made over at all #irony

3. Context in Social Media

The interaction between the different sources of contextual information discussed so far highlights a set of challenging issues in the semantics–pragmatics interface, not all

⁷ See McNally (2013) for an interesting discussion on that topic.

of which are solved and clear at the theoretical level. In addition, the NLP challenge is how to take these insights about different types of context and make good use of them in applications—in particular in applications that involve social media content. In this section, we review recent developments in processing social media language that incorporate the role of context.

3.1 On the Role of Discourse Phenomena

Discourse structure in social media conversations (like Twitter multilogues, i.e., conversations between users via the reply-to relation) differs in a number of aspects from that of “classical” dialogues (i.e., human–human and human–machine spoken dialogues). Indeed, some specific features such as Twitter @-mentions and hashtags may pose some problems regarding the choice of the appropriate unit of analysis (sentence, discourse unit, etc.) and level of the discourse structure these units should be embedded (Sidarenka, Bisping, and Stede 2015). In addition, social media corpora are composed of follow-up conversations, where topics are dynamic over conversation threads—that is, not necessarily known in advance. For example, posts on a forum or tweets are often responses to earlier posts, and the lack of context makes it difficult for machines to figure out, for example, whether the post is in agreement or disagreement.

Discourse contextual phenomena in social media can be leveraged in several ways, as discussed in the next sections.

3.1.1 Discourse Structure and Coherence Modeling. Although the analysis of discourse structure for traditionally written text is now well established (Lin, Kan, and Ng 2009; Hernault et al. 2010; Feng and Hirst 2014; Joty, Carenini, and Ng 2015), there is little work on applying discourse theories to social media texts. Among them, Sidarenka, Bisping, and Stede (2015) study how coherence is achieved in social media conversations relying on Rhetorical Structure Theory (Mann and Thompson 1988). They propose a scheme to manually annotate tweets according to Rhetorical Structure Theory principles and found that up to 40% of German tweets are part of conversations, and that answer-relations create discourse trees. The analysis of Twitter-specific phenomena reveals that URLs carry communicative content (such as Inform, Opening, Suggestion). Similarly, discourse relations (such as Elaboration, Exemplification, Evaluation) are rarely explicit (only 20% of the cases). They also observe that causal connectives are frequent in Twitter: 1.7% of the tweets and 2.6% of the replies.

Following the entity grid coherent model (Barzilay and Lapata 2008), Joty, Nguyen, and Mohiuddin (2018) also focus on the problem of coherence in asynchronous conversations. The authors propose a neural model to predict the underlying thread structure of fora conversations. The model has also been applied in reconstructing thread structures.

Finally, Perret et al. (2016) propose the first discourse parser for multi-party chat dialogues using integer linear programming. They investigate both treelike and non-treelike full discourse structures, achieving an F-measure of 0.531. These results are encouraging and open interesting future directions in discourse parsing of social media conversations.

3.1.2 Argumentation Mining. Specific argumentative discourse relations are of particular importance in social media. Indeed, a user often not only reports facts, expresses opinion, and engages with the reader, but also presents arguments in a certain order and with certain organization. These arguments are structured in terms of a set of premises

that provide the evidence or the reasons for or against a conclusion. Tracking arguments in text, also known as argumentation mining, consists of first identifying arguments (i.e., separating arguments from non-arguments), then their argumentative structure (including the premises, conclusion, and the connections between them such as the argument and counter-argument relationships). Argumentation mining in Twitter has been studied by Bosc, Cabrio, and Villata (2016), who propose a binary classifier for argument identification. Dusmanu, Cabrio, and Villata (2017) go further by separating personal opinions from actual facts, and detecting the source of such facts to allow for provenance verification.

Argumentation mining in social media has given rise to new tasks such as detecting agreements and disagreement in conversations (Allen, Carenini, and Ng 2014), counterfactual recognition (Son et al. 2017), identification of controversial topics (Addawood and Bashir 2016), stance/rumor detection (Zubiaga et al. 2016), and fact-checking (Baly et al. 2018). Argumentation and stancetaking are further discussed later in this special issue (cf. Cocarascu et al. and Kiesling et al., respectively).

3.1.3 Intention Detection. Another line of research concerns *intention* prediction.⁸ Analyzing intentions in conversations is an old topic in natural language understanding, where the goal is to detect what the speaker plans to pursue with their speech acts (Allen and Perrault 1980). Compared with the Web search community, where predicting user intentions from search queries and/or the user's click behavior has been extensively studied (Chen et al. 2002), there is little research that investigates how to extract intentions from users' free text.

The first attempt was the use of indirect speech acts to detect e-mails requesting actions (Cohen, Carvalho, and Mitchell 2004). E-mail intent detection is treated as a binary classification problem (request vs. nonrequest), leaving apart the difficult determination of the precise extent of the text that conveys this request. With the rise of social media, capturing intentions from user-generated content has become an emerging research topic. Most approaches aim at assigning predefined speech-act categories, like ASSERTION, RECOMMENDATION, REQUEST, QUESTION, COMMENT. Methods vary from supervised learning with bag-of-words representations to unsupervised models exploiting surface features (e.g., punctuations, emoticons), sentence-internal structure (e.g., parts of speech, dependency relations) (Zarisheva and Scheffler 2015; Vosoughi and Roy 2016), or to a little extent, the conversational dependencies between sentences, collapsing the set of user's writings (tweets) into the same sequence (Joty and Hoque 2016).

3.1.4 Conversational Thread and Topic as Key Contextual Factors. Discourse analysis of social media is a growing field of interest in linguistics in general and in discourse analysis in particular, with a significant amount of the research published in journals such as *Discourse Studies* or *Journal of Pragmatics* analyzing social media language, and even an entire journal devoted to this field (*Discourse, Context & Media*, published by Elsevier). Although the study of discourse and context in computational linguistics is perhaps not central, leveraging the context provided by the conversation thread and topic has recently been the center of many NLP applications. Perhaps the best example comes from sentiment analysis where conversations are used to enhance the performance of polarity detection. Indeed, although neighboring tweets tend to share similar polarity,

⁸ We use the term **intention** as a broader term that covers desires, plans, goals, and preferences.

the polarity orientation of the root (i.e., the original post/tweet) is usually shifted during the reply process (Huang, Cao, and Dong 2016). Vanzo, Croce, and Basili (2014) model polarity detection as a sequential classification task over streams of tweets about the same topic and observe an improvement of about 20% in F1 measure compared with approaches that do not account for the history of preceding posts. Ren et al. (2016) incorporate word embedding vectors extracted from both the current tweet's content and the conversation context into a neural network, and measure the role of context based on history tweets of the same author, which can serve as a prior for a tweet's sentiment. The context-based neural model gains more than 10% in macro F-measure.

Figurative language processing is another area of research where conversation plays a crucial role. With social media texts being very short, it is often difficult to recognize sarcasm or irony on the basis of the content of an utterance taken in isolation. Hence, the context provided by the preceding messages can help in detecting the incongruity between the literal meaning of an utterance and its intended meaning. Several approaches have been proposed to leverage such context, like Bamman and Smith (2015), who explore the properties of the author (e.g., profile information and historical salient terms), the audience (author/addressee topics), and the immediate communicative environment (previous tweets); and Wallace, Choe, and Charniak (2015), who exploit signals extracted from the conversational threads to which the comments belong. For a general discussion of context-based approaches to irony/sarcasm detection, we refer the reader to Joshi, Bhattacharyya, and Carman (2017).

Topic prediction can also benefit from document/posts sequential structure. For example, Ghosh et al. (2016) recently propose Contextual Long-Short Term Memory (CLSTM), a new sequence learning model that extends the recurrent neural network LSTM by incorporating contextual features. CLSTM has been used for sentence topic prediction: Given the words and the topic of the current sentence, predict the topic of the next sentence.

3.2 On the Role of Other Contextual Phenomena

In addition to the discursive contextual phenomena that are mainly driven from posts' conversation structure, there are many other types of context that can be combined with linguistic content. Among them, we focus now on demographic information and social network structure.

3.2.1 Demographic Information. This refers to author-related information like age, gender, race, income, location, political orientation, and other demographic categories. Two lines of research have recently gained relevance in the NLP community to derive demographic information from texts: author profiling and author identification (Rosso et al. 2018; Stamatatos et al. 2018). In the first task, information such as the author's age and gender can be predicted, as authors who share similar demographic traits also share similar linguistic patterns. In the second task, given a group of potential authors, the goal is to determine the right one (also known as authorship attribution). Whereas most approaches mainly rely on lexical features derived from the linguistic content of the message alone, recent approaches propose to account for discourse structure (Wanner and Soler 2017).

When available, author-related information has been extensively used in different NLP tasks, including sentiment/emotion analysis. For instance, several studies have found strong correlations between the expression of subjectivity and gender (for example, some subjective words will be used by men, but never by women, and vice versa),

and leverage these correlations for gender identification (Burger et al. 2011; Volkova and Bachrach 2016). Stylometric and personality features of users have also been used for sarcasm detection (Hazarika et al. 2018).

Detecting the location of the social media users provides another type of demographic information useful in various applications. This information can be directly available from user profiles or other meta-data (such as GPS information for posted messages). When it is not available, it can be predicted based on the network structure (“you are where your friends are”) or relations between those who follow and those who are followed (Rout et al. 2013) or based on the content of the posted messages. The latter content-based approaches extract information about the use of language, the main topics discussed, the named entities mentioned frequently, and so on. (Eisenstein et al. 2010; Han, Cook, and Baldwin 2012; Liu and Inkpen 2015). The accuracy of these methods is not high, but it can be improved by combining content-based approaches with the contextual information provided by the network structure and other location-indicative meta-data.

3.2.2 Social Network Structure. In social media, social relationships between users enable grouping users into specific communities. A community is often not identified in advance, but its users are expected to share common goals: circles of friends, members, groups of topically related conversations, and so forth. Drawing from the assumption that users connected in the social network (e.g., via followers, mentions, reply-to) or that belong to the same community may have similar subjective orientations, several studies show that users’ social relationships can enhance sentiment analysis (Tan et al. 2011). For example, Huang, Singh, and Atrey (2014) showed that modeling the social network structure improves accuracy when detecting cyber-bullying messages.

4. Overview of the Articles in this Special Issue

This issue aimed to study how the treatment of linguistic phenomena, in particular at the discourse level, can benefit NLP-based social media systems, and help such systems advance beyond representations that include only bags of words or bags of sentences. Discourse and pragmatic information can also help move beyond sentence-level approaches that typically account for local contextual phenomena relying on dedicated lexicons and shallow or deep syntactic parsing. More importantly, the aim of this issue is to show that incorporating linguistic insights, discourse information, and other contextual phenomena, in combination with the statistical exploitation of data, can result in an improvement over approaches that take advantage of only one of those perspectives.

We received a total of 15 submissions, reflecting a significant interest in these phenomena in the computational linguistics community. After a rigorous review process, we selected six articles, covering various aspects of the topic. The selected articles address deep issues in linguistics, computational linguistics, and social science. The special issue is structured around three main themes, according to the type of context considered in each article:

- *Social context:* The focus here is on the social and relational meaning in online conversations from a theoretical point of view (Kiesling et al.).
- *Conversation turns and common-sense knowledge:* Here, we group papers that study phenomena for which people make inferences in their everyday use

of language, focusing on inferences that are drawn when searching for the figurative meaning of an utterance (Ghosh et al.; Van Hee et al.).

- *Conversational context*: The third part focuses on the role of discourse phenomena in processing social media conversations, including topicality (Li et al.), speech acts (Joty and Mohiuddin), and argumentation (Cocarascu and Toni).

The rest of this section provides a brief introduction to each of the six accepted papers.

The article by Kiesling et al. (“Interactional Stancetaking in Online Forums”) investigates thread structure and linguistic properties of stancetaking from the online platform Reddit. Stancetaking captures the speaker’s (or writer’s) relationship to the topic of discussion, the interlocutor, or audience, and the talk (or writing) itself. The authors first propose a new data set where conversation threads are annotated according to three linked stance dimensions: affect, investment, and alignment. These dimensions are then predicted relying on lexical features. The quantitative and qualitative results of this study show that stance utterances tend to pattern in coherent conversational threads.

Li et al. (“A Joint Model of Conversational Discourses and Latent Topics on Microblogs”) extract topics from microblog messages, a challenging task given the data sparsity in short messages that often lack structure and context. To address this issue, the authors represent microblog messages as conversation trees based on their reposting and replying relations, and propose an unsupervised model that jointly learns word distributions to identify the different functions of conversational discourse and various latent topics to represent content-specific information embedded in microblog messages. Their experiments show that the proposed joint model on topic coherence outperform state-of-the-art models. The output from the joint model is then used for microblog summarization: By additionally capturing word distributions for different sentiment polarities, the jointly modeled discourse and topic representations can effectively indicate summary-worthy content in microblog conversations.

The article by Ghosh et al. (“Sarcasm Analysis Using Conversation Context”) studies the role of conversation to detect sarcasm in tweets and discussion forums. The context considered here concerns the current turn as well as the prior and the succeeding one (when available). In order to show to what extent modeling of conversation context helps in sarcasm detection, the authors investigate both classical learning models with linguistically motivated discrete features and several types of LSTM networks (conditional LSTM network, LSTM networks with sentence-level attention). The models were tested on different corpus genre data sets and the results show that attention models achieve significant improvement when using the prior turn as context for all the data sets. To better measure the difficulty of the task, the authors perform a qualitative analysis of attention weights produced by the LSTM models and discuss the results compared with human performance on the task.

In the article by Van Hee et al. (“We Usually Don’t Like Going to the Dentist: Using Common Sense to Detect Irony on Twitter”), the role of context in figurative language detection is also explored. Compared with Ghosh et al., who focus on conversational context, Van Hee et al. target common sense and connotative knowledge and propose to model implicit or prototypical sentiment (e.g., “flight delays,” “going to the dentist” generally convey negative sentiment) in the framework of automatic irony detection in tweets. Their approach uses a support vector machine classifier relying on lexical,

syntactic, and semantic features, with a particular focus on lexical and semantic features that have been extended with language model features and word cluster information. The results show that applying sentiment analysis using SenticNet and real-time crawled tweets is a viable method to determine the implicit sentiment related to that concept or situation.

Cocarascu and Toni (“Combining Deep Learning and Argumentative Reasoning for the Analysis of Social Media Textual Content Using Small Data Sets”) propose a method to check whether news headlines support statements from tweets, to allow for fact-checking. Their deep learning method extracts argumentative relations of attack and support. Then they use the proposed method to extract bipolar argumentation frameworks from reviews, to help detect whether they are deceptive. They show experimentally that the method performs well in both settings. In particular, in the case of deception detection, the method contributes a novel argumentative feature that, when used in combination with other features in standard supervised classifiers, outperforms the latter even on small data sets.

The last article in this special issue, by Joty and Mohiuddin (“Modeling Speech Acts in Asynchronous Conversations: A Neural-CRF Approach”), presents a method for speech act recognition, a problem that has long been a concern in the spoken dialogue research community, and one that poses particular problems in online social media communication, which tends to be asynchronous. Joty and Mohiuddin train LSTM-RNNs using conversational word embeddings. This is a significant result, as they show that word embeddings trained on a related domain improve the performance of the system. The contribution of this article is to incorporate context in the form of dependencies across sentences. It is clear from the literature that conversation structure is relevant when interpreting speech acts. The authors propose to model it as a graph structure, given the nonlinear nature of asynchronous conversation. In addition, Joty and Mohiuddin work from the hypothesis that, when representing sentence meaning, word order is important, and should be preserved. Although this does not seem like a revolutionary concept, word order is often disregarded in “classic” machine learning approaches, and in modern vector representations of text.

5. Conclusions and Future Directions

We hope that this special issue contributes to a deeper understanding of the role of different types of context and their interaction to process social media data from the perspective of discourse interpretation. We believe that we are entering a new age of mining social media data, one that extracts information not just from individual words, phrases, and tags, but also uses information from discourse and the wider context. Most of the “big data” revolution in social media analysis has examined words in isolation—a bag-of-words approach. We believe it is possible to investigate big data, and social media data in general, by exploiting contextual information.

To achieve that purpose, we need to first develop tools to automatically determine the structure of discourse, including discourse relations, argumentation, and threads in conversations such as those found in Twitter and other social media. This is an interdisciplinary enterprise that needs to address deep issues in both linguistics and computational linguistics, including the analysis of the discursive properties of social media content and the empirical study of how these properties are deployed in different corpus genres through corpus annotation. We need to propose new solutions in various use cases including sentiment analysis, detection of offensive content, and intention

detection. These solutions need to be reliable enough in order to prove their effectiveness against shallow bag of words approaches.

Another direction of research that we encourage is to further explore the interactions between content and extra-linguistic or extra-textual features, in particular time, place, author profiles, demographic information, conversation thread, and network structure.

Acknowledgments

We would like to thank all the authors who submitted articles and all the reviewers for their time and effort. We also greatly thank the journal editors, Paola Merlo and Hwee Tou Ng, for their guidance and support during the entire process.

References

- Addaood, Aseel and Masooda Bashir. 2016. "What is your evidence?" A study of controversial topics on social media. In *Proceedings of the Third Workshop on Argument Mining*, ArgMining 2016, pages 1–11, Berlin, Germany.
- Aiello, Luca Maria, Georgios Petkos, Carlos J. Martín, David Corney, Symeon Papadopoulos, Ryan Skraba, Ayse Göker, Ioannis Kompatsiaris, and Alejandro Jaimes. 2013. Sensing trending topics in Twitter. *IEEE Transaction of Multimedia*, 15(6):1268–1282.
- Allen, J. F. and C. R. Perrault. 1980. Analyzing intention in utterances. *Artificial Intelligence*, 15(3):143–178.
- Allen, Kelsey, Giuseppe Carenini, and Raymond T. Ng. 2014. Detecting disagreement in conversations using pseudo-monologic rhetorical structure. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP 2014, pages 1169–1180, Doha.
- Asher, Nicholas and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press.
- Attardo, Salvatore. 2000. Irony as relevant inappropriateness. *Journal of Pragmatics*, 32(6):793–826.
- Austin, John Langshaw. 1962. *How to Do Things with Words*. Oxford.
- Bach, Kent. 1997. The semantics-pragmatics distinction: What it is and why it matters. VS Verlag für Sozialwissenschaften. pages 33–50.
- Baly, Ramy, Mitra Mohtarami, James R. Glass, Lluís Màrquez, Alessandro Moschitti, and Preslav Nakov. 2018. Integrating stance detection and fact checking in a unified corpus. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 21–27, New Orleans, LA.
- Bamman, David and Noah A. Smith. 2015. Contextualized sarcasm detection on Twitter. In *Proceedings of the International Conference on Web and Social Media*, ICWSM 2015, pages 574–577, Oxford, UK.
- Barzilay, Regina and Mirella Lapata. 2008. Modeling local coherence: An entity-based approach. *Computational Linguistics*, 34(1):1–34.
- Benamara, Farah, Nicholas Asher, Yannick Mathieu, Vladimir Popescu, and Baptiste Chardon. 2016. Evaluation in Discourse: a Corpus-Based Study. *Dialogue and Discourse*, 7(1):1–49.
- Benamara, Farah, Maite Taboada, and Yannick Mathieu. 2017. Evaluative Language Beyond Bags of Words: Linguistic Insights and Computational Applications. *Computational Linguistics*, 43(1):201–264.
- Bos, Johan. 2011. A survey of computational semantics: Representation, inference and knowledge in wide-coverage text understanding. *Language and Linguistics Compass*, 5(6):336–366.
- Bosc, Tom, Elena Cabrio, and Serena Villata. 2016. Tweeties squabbling: Positive and negative results in applying argument mining on social media. In *Proceedings of Computational Models of Argument*, COMMA 2016, pages 21–32, Potsdam.
- Bunt, Harry. 2001. From lexical item to discourse meaning: Computational and representational tools. In *Computing Meaning, volume 77 of Studies in Linguistics and Philosophy*. Springer Netherlands, pages 1–10.
- Bunt, Harry and Bill Black. 2000. *The ABC of computational pragmatics*. John Benjamins, pages 1–46.
- Burger, John D., John Henderson, George Kim, and Guido Zarrella. 2011. Discriminating gender on Twitter. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1301–1309, Edinburgh.

- Chen, Zheng, Fan Lin, Huan Liu, Yin Liu, Wei-Ying Ma, and Liu Wenyan. 2002. User intention modeling in Web applications using data mining. *World Wide Web*, 5(3):181–191.
- Cohen, William W., Vitor R. Carvalho, and Tom M. Mitchell. 2004. Learning to classify email into “speech acts.” In Dekang Lin and Dekai Wu, editors, *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP 2004, pages 309–316, Barcelona.
- Deitrick, William and Wei Hu. 2013. Mutually enhancing community detection and sentiment analysis on twitter networks. *Journal of Data Analysis and Information Processing*, 1(3):19–29.
- Dusmanu, Mihai, Elena Cabrio, and Serena Villata. 2017. Argument mining on Twitter: Arguments, facts and sources. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, EMNLP 2017, pages 2317–2322, Copenhagen, Denmark.
- Eisenstein, Jacob, Brendan O’Connor, Noah A. Smith, and Eric P. Xing. 2010. A latent variable model for geographic lexical variation. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1277–1287, Cambridge, MA.
- Farzindar, Atefeh and Diana Inkpen. 2017. *Natural Language Processing for Social Media*. Morgan & Claypool Publishers.
- Feng, Vanessa Wei and Graeme Hirst. 2014. A linear-time bottom-up discourse parser with constraints and post-editing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 511–521, Baltimore, MD.
- Ghosh, Aniruddha, Guofu Li, Tony Veale, Paolo Rosso, Ekaterina Shutova, John A. Barnden, and Antonio Reyes. 2015. Semeval-2015 task 11: Sentiment analysis of figurative language in Twitter. In *Proceedings of the 9th International Workshop on Semantic Evaluation*, SemEval@NAACL-HLT 2015, pages 470–478, Denver, CO.
- Ghosh, Shalini, Oriol Vinyals, Brian Strope, Scott Roy, Tom Dean, and Larry P. Heck. 2016. Contextual LSTM (CLSTM) models for large scale NLP tasks. *CoRR*, abs/1602.06291.
- Grice, H. Paul. 1975. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Speech Acts. Syntax and Semantics, Volume 3*, Academic Press, pages 41–58.
- Grosz, B. J., Aravind K. Joshi, and Scott Weinstein. 1995. Centering: A framework for modelling the local coherence of discourse. *Computational Linguistics*, 21(2):203–225.
- Grosz, Barbara J. and Candace L. Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.
- Halliday, Alexander Kirkwood and Ruqaiya Hasan. 1976. *Cohesion in English*. Routledge.
- Han, Bo, Paul Cook, and Timothy Baldwin. 2012. Geolocation prediction in social media data by finding location indicative words. In *Proceedings of COLING 2012*, pages 1045–1062, Mumbai.
- Hazarika, Devamanyu, Soujanya Poria, Sruthi Gorantla, Erik Cambria, Roger Zimmermann, and Rada Mihalcea. 2018. CASCADE: Contextual sarcasm detection in online discussion forums. In *Proceedings of the 27th International Conference on Computational Linguistics*, ACL 2018, pages 1837–1848, Santa Fe, NM.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts.
- Hernault, H., H. Prendinger, D. duVerle, and M. Ishizuka. 2010. Hilda: A discourse parser using support vector machine classification. *Dialogue and Discourse*, 1(3):1–33.
- Hobbs, Jerry. 1979. Coherence and coreference. *Cognitive Science*, 3(8):67–90.
- Huang, Minlie, Yujie Cao, and Chao Dong. 2016. Modeling rich contexts for sentiment classification with LSTM. *CoRR*, abs/1605.01478.
- Huang, Qianjia, Vivek Kumar Singh, and Pradeep Kumar Atrey. 2014. Cyber bullying detection using social and textual analysis. In *Proceedings of the 3rd International Workshop on Socially-Aware Multimedia*, SAM ’14, pages 3–6, New York, NY.
- Inkpen, Diana, Ji Liu, Atefeh Farzindar, Farzaneh Kazemi, and Diman Ghazi. 2015. Detecting and disambiguating locations mentioned in Twitter messages. In *Computational Linguistics and Intelligent Text Processing*, CICLing, pages 321–332, Cairo.
- Janssen, Theo M. V. 2001. Frege, contextuality and compositionality. *Journal of Logic, Language and Information*, 10(1):115–136.
- Jaszczolt, K. M. 2012. Semantics and pragmatics: The boundary issue. In K. von Stechow, P. Portner, and C. Maienborn, editors, *Semantics: An International Handbook of Natural Language Meaning*,

- Mouton de Gruyter, Berlin, pages 306–332.
- Joshi, Aditya, Pushpak Bhattacharyya, and Mark J. Carman. 2017. Automatic sarcasm detection: A survey. *ACM Computing Surveys*, 50(5):1–22.
- Joty, Shafiq, Giuseppe Carenini, and Raymond Ng. 2015. CODRA: A novel discriminative framework for rhetorical analysis. *Computational Linguistics*, 41(3):385–435.
- Joty, Shafiq R. and Enamul Hoque. 2016. Speech act modeling of written asynchronous conversations with task-specific embeddings and conditional structured models. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1746–1756, Berlin.
- Joty, Shafiq R., Dat Tien Nguyen, and Muhammad Tasnim Mohiuddin. 2018. Coherence modeling of asynchronous conversations: A neural entity grid approach. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, ACL 2018, pages 558–568, Melbourne.
- Kamp, Hans and Uwe Reyle. 1993. *From Discourse to Logic*. Dordrecht.
- Karoui, Jihen, Farah Benamara, Véronique Moriceau, Nathalie Aussenac-Gilles, and Lamia Hadrich-Belguith. 2015. Towards a contextual pragmatic model to detect irony in tweets. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 644–650, Beijing, China.
- Korta, Kapa and John Perry. 2015. Pragmatics. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2015/entries/pragmatics/>.
- Krifka, Manfred. 2002. Embedded speech acts. In *Proceedings of the Workshop In the Mood*, Frankfurt.
- Lenci, Alessandro. 2006. The lexicon and the boundaries of compositionality. *Acta Philosophica Fennica*, 78:303–320.
- Lenci, Alessandro. 2018. Distributional models of word meaning. *Annual Review of Linguistics*, 4(1):151–171.
- Lin, Ziheng, Min-Yen Kan, and Hwee Tou Ng. 2009. Recognizing implicit discourse relations in the Penn discourse treebank. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 343–351, Singapore.
- Liu, Ji and Diana Inkpen. 2015. Estimating user location in social media with stacked denoising auto-encoders. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 201–210, Denver, CO.
- Londhe, Nikhil, Rohini K. Srihari, and Vishrawas Gopalakrishnan. 2016. Time-independent and language-independent extraction of multiword expressions from Twitter. In *26th International Conference on Computational Linguistics, COLING*, pages 2269–2278, Osaka.
- Mann, William C. and Sandra A. Thompson. 1988. Rhetorical Structure Theory: Toward a functional theory of text organization. *Text*, 8(3):243–281.
- McNally, Louise. 2013. Semantics and pragmatics. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4:285–297.
- Montague, Richard. 1974. English as a formal language. In Richmond H. Thomason, editor, *Formal Philosophy: Selected Papers of Richard Montague*, Yale University Press, New Haven, CT, pages 188–222.
- Mukherjee, Subhabrata and Pushpak Bhattacharyya. 2012. Sentiment analysis in Twitter with lightweight discourse analysis. In *Proceedings of International Conference on Computational Linguistics, COLING 2012*, pages 1847–1864, Mumbai.
- Perret, Jérémy, Stergos D. Afantenos, Nicholas Asher, and Mathieu Morey. 2016. Integer linear programming for discourse parsing. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 99–109, San Diego, CA.
- Persing, Isaac and Vincent Ng. 2014. Vote prediction on comments in social polls. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1127–1138, Doha.
- Polanyi, Livia and Martin van den Berg. 2011. Discourse structure and sentiment. In *Data Mining Workshops (ICDMW)*, pages 97–102, Vancouver.
- Prasad, Rashmi, Bonnie Webber, and Aravind Joshi. 2014. Reflections on the Penn Discourse Treebank, comparable corpora, and complementary annotation. *Computational Linguistics*, 40(4):921–950.
- Pustejovsky, James. 1995. *The Generative Lexicon*. MIT Press.
- Recanati, François. 2004. *Literal Meaning*. Literal Meaning. Cambridge University Press.

- Recanati, François. 2008. *Pragmatics and Semantics*. Blackwell Publishing LTD. pages 442–462.
- Ren, Yafeng, Yue Zhang, Meishan Zhang, and Donghong Ji. 2016. Context-sensitive twitter sentiment classification using neural network. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI 2016, pages 215–221, Phoenix, AZ.
- Rosso, Paolo, Francisco M. Rangel Pardo, Irazú Hernandez-Farias, Leticia C. Cagnina, Wajdi Zaghouni, and Anis Charfi. 2018. A survey on author profiling, deception, and irony detection for the Arabic language. *Language and Linguistics Compass*, 12(4):1–20.
- Rout, Dominic, Kalina Bontcheva, Daniel Preotiuc-Pietro, and Trevor Cohn. 2013. Where's @wally?: A classification approach to geolocating users based on their social ties. In *HyperText and Social Media 2013*, pages 11–20, Paris.
- Saloot, Mohammad Arshi, Norisma Idris, AiTi Aw, and Dirk Thorleuchter. 2016. Twitter corpus creation: The case of a Malay chat-style-text corpus (MCC). *Digital Scholarship in the Humanities*, 31(2):227–243.
- Searle, John R. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press.
- Sidarenka, Uladzimir, Matthias Bisping, and Manfred Stede. 2015. Applying Rhetorical Structure Theory to Twitter conversations. In *Proceedings of the Workshop on Identification and Annotation of Discourse Relations in Spoken Language (DiSpol)*, pages 1–2, Saarbrücken.
- Socher, Richard, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013*, pages 1631–1642, Seattle, WA.
- Son, Youngseo, Anneke Buffone, Joe Raso, Allegra Larche, Anthony Janocko, Kevin Zembroski, H. Andrew Schwartz, and Lyle Ungar. 2017. Recognizing counterfactual thinking in social media texts. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, pages 654–658, Vancouver.
- Sperber, Dan and Deirdre Wilson. 1981. Irony and the use-mention distinction. *Radical Pragmatics*, 49:295–318.
- Stamatatos, Efsthios, Francisco M. Rangel Pardo, Michael Tschuggnall, Benno Stein, Mike Kestemont, Paolo Rosso, and Martin Potthast. 2018. Overview of PAN 2018 - author identification, author profiling, and author obfuscation. In *CLEF 2018*, volume 11018 of *Lecture Notes in Computer Science*, pages 267–285, Springer.
- Taboada, Maite and William C. Mann. 2006. Rhetorical structure theory: Looking back and moving ahead. *Discourse Studies*, 8(3):423–459.
- Tan, Chenhao, Lillian Lee, Jie Tang, Long Jiang, Ming Zhou, and Ping Li. 2011. User-level sentiment analysis incorporating social networks. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining, SIGKDD*, pages 1397–1405, San Diego, CA.
- Tarski, Alfred. 1983. *Logic, semantics, metamathematics: Papers from 1923 to 1938*. Hackett Publishing Company, Indianapolis, IN.
- Turney, Peter D. and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligent Research*, 37(1):141–188.
- Utsumi, Akira. 1996. A unified theory of irony and its computational formalization. In *Proceedings of the International Conference on Computational Linguistics, COLING*, pages 962–967, Copenhagen.
- Vanzo, Andrea, Danilo Croce, and Roberto Basili. 2014. A context-based model for sentiment analysis in Twitter. In *Proceedings of the 25th International Conference on Computational Linguistics, COLING 2014*, pages 2345–2354, Dublin.
- Volkova, Svitlana and Yoram Bachrach. 2016. Inferring perceived demographics from user emotional tone and user-environment emotional contrast. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016*, pages 1567–1578, Berlin.
- Volkova, Svitlana, Glen Coppersmith, and Benjamin Van Durme. 2014. Inferring user political preferences from streaming communications. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014*, pages 186–196, Baltimore, MD.
- Vosoughi, Soroush and Deb Roy. 2016. Tweet acts: A speech act classifier for Twitter. In *Proceedings of International AAAI Conference on Web and Social Media, ICWSM 2016*, pages 711–715, Cologne.

- Wallace, Byron C., Do Kook Choe, and Eugene Charniak. 2015. Sparse, contextually informed models for irony detection: Exploiting user communities, entities and sentiment. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing, ACL-IJCNLP*, pages 1035–1044, Beijing.
- Wanner, Leo and Juan Soler. 2017. On the relevance of syntactic and discourse features for author profiling and identification. In *EACL 2017*, pages 681–687, Valencia.
- Webber, Bonnie, Markus Egg, and Valia Kordoni. 2012. Discourse structure and language technology. *Natural Language Engineering*, 18(4):437–490.
- West, Robert, Hristo S. Paskov, Jure Leskovec, and Christopher Potts. 2014. Exploiting social network structure for person-to-person sentiment analysis. *Transactions of the Association of Computational Linguistics (TACL)*, 2:297–310.
- Zarisheva, Elina and Tatjana Scheffler. 2015. Dialog act annotation for Twitter conversations. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, SIGDIAL 2017*, pages 114–123, Prague.
- Zimmermann, T. E. 2013. The Oxford handbook of compositionality. In Wolfram Hinzen, Edouard Machery and Markus Werning, editors, *Compositionality Problems and How to Solve Them*, Oxford University Press, pages 81–106.
- Zubiaga, Arkaitz, Elena Kochkina, Maria Liakata, Rob Procter, and Michal Lukasik. 2016. Stance classification in rumours as a sequential task exploiting the tree structure of social media conversations. In *Proceedings of the 26th International Conference on Computational Linguistics: Technical Papers, COLING 2016*, pages 2438–2448, Osaka.

