May 22nd, 9:00 AM - May 25th, 5:00 PM

# Integrating natural language processing and pragmatic argumentation theories for argumentation support

Mark Aaknus
*Rutgers, The State University of New Jersey, School of Communications & Information*

Smaranda Muresan
*Rutgers, The State University of New Jersey, School of Communications & Information*

Nina Wacholder
*Rutgers, The State University of New Jersey, School of Communications & Information*

# Integrating natural language processing and pragmatic argumentation theories for argumentation support

## MARK AAKHUS

*Communication*
*Laboratory for the Study of Applied Language Technology and Society (SALTS)*
*School of Communication & Information*
*Rutgers, The State University of New Jersey*
*4 Huntington St., New Brunswick, NJ 08901 USA*
*aakhus@rutgers.edu*

## SMARANDA MURESAN

*Library and Information Science*
*Laboratory for the Study of Applied Language Technology and Society (SALTS)*
*School of Communication & Information*
*Rutgers, The State University of New Jersey*
*4 Huntington St., New Brunswick, NJ 08901 USA*
*smuresan@rutgers.edu*

## NINA WACHOLDER

*Library and Information Science*
*Laboratory for the Study of Applied Language Technology and Society (SALTS)*
*School of Communication & Information*
*Rutgers, The State University of New Jersey*
*4 Huntington St., New Brunswick, NJ 08901 USA*
*ninwac@rutgers.edu*

ABSTRACT: Natural language processing (NLP) research and design that aims to model and detect opposition in text for the purpose of opinion classification, sentiment analysis, and meeting tracking, generally excludes the interactional, pragmatic aspects of online text. We propose that a promising direction for NLP is to incorporate the insights of pragmatic, dialectical theories of argumentation to more fully exploit the potential of NLP to offer sound, robust systems for various kinds of argumentation support.

KEYWORDS: argumentation, argument ontology, collective intelligence, computational tools for argument support, computer supported argument visualization, disagreement space, epistemic capabilities, natural language processing, sentiment identification

## 1. THE PROSPECTS FOR COMPUTATIONAL ARGUMENTATION SUPPORT

A long-standing motivation for computation has been the development of socially intelligent systems that augment human reasoning and interaction (e.g., Bush, 1945; Englebardt, 1962; Licklider, 1960). A fundamental challenge, however, lies in

Mohammed, D., & Lewiński, M. (Eds.). *Virtues of Argumentation. Proceedings of the 10th International Conference of the Ontario Society for the Study of Argumentation (OSSA), 22-26 May 2013.* Windsor, ON: OSSA, pp. 1-12.

developing methods that go beyond aggregating disparate pieces of information toward methods for understanding the collective intelligence produced when interacting collectives engage in making-sense of prudent courses of action relative to some social, political, economic, medical, or environmental matter. Such a method might enable the articulation of what is more arguable within an interacting collective and what is less arguable and thus afford modeling the epistemic capabilities of an interacting collective by discovering how the collective manages disagreement. Such a method would help collectives, participants and observers detect and track lines of disagreement in a discussion, and articulate sources of contestation and the manner in which matters are made contestable, while bringing to the surface the common-places and lines of reasoning interacting parties use to oppose and to construct arguments.

There are two basic computational approaches for representing the argumentative aspects of messages exchanged online that address this challenge: one is Natural Language Processing (NLP) and the other is Computer Supported Argument Visualization (CSAV). NLP has made great strides in identifying sentiment and opinion but does not yet provide the deep semantic and pragmatic analysis necessary for understanding and supporting large-scale argumentation by communities. CSAVs have provided some rich ontologies for representing the argumentative relations among contributions made by interacting collectives but has ultimately been constrained by scale. But even when the strengths of both are combined, extant approaches grounded in NLP and CSAV remain limited in articulating the reasoning of interacting communities.

Two interrelated problems underlying extant approaches are discussed here to outline two requirements for designing computational support for argumentation. NLP usually does not model argumentation in terms of a response-centered approach. CSAV focuses on designing ontologies of argument relations rather than on the potential for computation to be part of a method for discovering the practical ontologies communities employ when managing disagreement and constructing issues.

## 2. AUTOMATIC TEXT CLASSIFICATION AND ARGUMENTATION

Within Natural Language Processing (NLP) there has been a significant amount of work on identifying sentiment and opinions in text. Automatic text classification makes it possible to represent the content (e.g., opinions) of what is being said, the stance from which it is being developed (e.g., sentiments), and the location of differences of opinion (e.g., dialogue zones). Most state-of-the-art NLP systems use machine learning to assigns labels (classes) to segments of text by taking advantage of known properties of language associated with an action such as expression of an opinion or disagreement with someone else's opinion. Besides counting frequency of words, phrases and grammatical parts of speech, shallow linguistic analysis can, for example, find named entities such as people, places and times and identify the main verb in a sentence along with its syntactic and semantic roles. To assign a label to a text segment or to identify the relationship between two segments, supervised machine learning is frequently used (Sebastiani, 2002). In the initial stage of supervised machine learning, the input is 'training text' that has been accurately labeled by classes. For sentiment classification, the machine learning system might use text that has been labeled as 'positive sentiment', 'negative sentiment' or 'neutral', as its training data. Machine learning uses sophisticated computational techniques to build statistical models of the distribution of features such as word frequency in the training text and to identify the combination of features that support the most accurate classification of text segments. The output of the initial (training) stage is a classifier model. This model is then used to assign classes to a test set, that is new text that is previously 'unseen' by the computer system. Classification accuracy of about 90% is typically expected for systems that will be used in real world applications. However, the labels that NLP assigns are localized and static and are based primarily on shallow linguistic analysis. This falls far short of deep semantic and pragmatic analysis required to model or support the reasoning of interacting collectives.

### *2.1 Approaches to classification of text segments as argumentative with NLP*

Automatic classification is a key component of NLP systems used to identify segments of text that express opinion and sentiment. Some work on sentiment identification simply assigns text to one of two classes. For example, Pang, Lee, & Vaithyanathan (2002) classify movie reviews as positive (thumbs up) or negative (thumbs down). Beyond simple binary classification of textual units, more granular classifications have been developed to identify opinion holders (e.g., Kim and Hovy, 2006) and assess strength of opinion (e.g., Wilson, Wiebe, & Hwa, 2004). There is also work focused on extracting opinion sentences (e.g., Hu & Liu, 2004; Popescu & Etzioni, 2005) and on identifying reasons for opinions (e.g., Kim & Hovy, 2006). A technique for improving the classification of opinions and sentiment involves the representation of content using richer features such as dialogue context and discourse, in conjunction with lexical features (e.g., Galley, McKeown, Hirschberg, & Shriberg, 2004; Wilson, Wiebe, & Hoffmann, 2005; Agarwal, Biadsy, & McKeown, 2009). Such approaches can provide superior results in classifying units of text as agreement or disagreement, as compared to just lexical and phrasal features (e.g., Hillard, Ostendorf, & Shriberg, 2003; Somasundaran, Namata, & Getoor, 2009; Thomas, Pang, & Lee, 2006).

Discourse features can also be used to label the beginning of a new turn in a transcript (Hawes, Lin, & Resnick, 2009). Dialogue approaches rely on identifying the role dialogue acts play in formulating decisions using automatic classification of text (Biu & Peters, 2010; Hsueh & Moore, 2007). Such techniques suggest that text can be classified as having zones, such as zones of conflict and cooperation, including locating where the discussion of action items occurs (Bunt, Alexandersson, Carletta, Choe, Chengyu Fang, Hasida, … Traum, 2010; Pallotta & Delmonte 2011). Classifying sentences in terms of participant, relation, and entity can show how the sentence plays a role in a planning dialogue (Carenini & Murray, 2009; Pallota, Niekrasz, & Purver, 2005; Pallotta, Seretan, & Ailomaa, 2007).

Despite the impressive progress in identification of sentiment, opinion and zones of differences, much work is needed to move beyond the simplistic representation of whole sentences or even entire documents as a single class when spans of text are often multi-functional. A key limiting factor has been the reliance on shallow linguistic approaches that aim to represent target text spans but do so in ways that do not adequately capture nuance or complex semantic and discourse relations. There are important parallels between the successes and limitations of extant NLP research and research on argumentative indicators that we acknowledge but do not examine in detail here. Although NLP approaches attend to textual clues, such as argumentative indicators, extant NLP approaches essentially ignore the sequence of discussion and the network of assumptions and presumptions available in the collective's discourse but mostly implicit in the text.

### *2.2 Design requirement: Identifying relationships between text segments with response centered analysis*

A key characteristic of argumentative discourse is that it unfolds sequentially but depends on the network of overarching presumptions and underlying assumptions. Actors participate in the sequential unfolding by raising doubts, suggesting commonalities or pointing out what is disagreeable and agreeable. Classically, this characteristic has been described in terms of the logic of the topoi where what is taken to be common-place and commonly held values can be used in generating doubts and disagreement(s). ”Disagreement Space” is a more contemporary account of this phenomenon that articulates the dynamic relationship between the explicit sequence of interaction and the tacit network of assumptions and presumptions at play when collectives are engaged in some activity (van Eemeren, Grootendorst, Jackson, & Jacobs, 1993). Disagreement can arise at any point when one or more actors engage in calling-out and making problematic some aspect of another actor's prior contribution for what it (could have) said or meant (van Eemeren et al., 1993). The argumentative relationships among contributions to a discussion are indicated through what is targeted and how it is called-out. Argumentative relations are constructed around the possible questions that can be raised about explicit and implicit matters, whether intended or unintended, and can be about, among other things, premise-conclusion relations, relevance of a contribution to an issue, pragmatic commitments of obligations and rights, and the relevance of a contribution to an activity.

Disagreement Space highlights how argumentative interaction is response centered in that argument is found in the way subsequent contributions relate to prior contributions. Shallow linguistic processing, however, does not treat language as a discourse that unfolds sequentially in time over turns while drawing on and developing underlying assumptions and overarching presumptions. Even when this complexity is acknowledged and treated with sophisticated machine learning techniques (e.g., Hawes et al., 2009), conventional NLP still does not model the patterns of relevance through which discourse holds together over time and across place. In part, this problem arises from limitations of state-of-the art computational models, but in part the problem lies in conceptualizing the meaning of a text for the purposes of computational analysis. An important alternative approach is inspired by moving from a shallow linguistic representation of text to a graph-based representation that attempts to encode utterance meaning in terms of the relation between text segments.

In order to achieve this desideratum, machine learning systems need to model richer, linguistically-inspired formalisms of language. One such formalism is the Lexicalized Well-Founded Grammar (LWFG), which combines syntax and semantics and is learnable from data (Muresan & Rambow, 2007; Muresan, 2011). Currently, LWFGs condition the interpretation of an utterance on linguistic context (``surrounding linguistic material" (Bunt, 1999)) and semantic/domain context (facts and knowledge of the domain encoded in an ontology). Once a LWFG grammar is learned, a LWFG parser and semantic/pragmatic interpreter map text to its underlying meaning representation encoded as a direct acyclic graph (DAG) (Muresan, 2008; Muresan, 2013). Vertices represent either concepts or instances of concepts expressed in nouns, verbs, adjectives, adverbs, and pronouns, or values of extra-ontological properties such as tense (e.g., present and future). Edges represent either semantic roles given by verbs, prepositions, adjectives and adverbs, or extra-ontological meaning properties such as tense, aspect, modality and negation. This meaning-level representation abstracts away from the surface form of the text and supports tractable inferences for extracting explicit and implicit information.

Modeling argumentation as response centered could rely on the LWFG formalism by exploiting two of its key features: the use of ontologies and graph-based meaning representations. In this approach, richer ontologies about argument can be used and the graph-based representation can be extended such that vertices are entire text segments and edges are argumentative relations. A segment of text is treated not as a set of words but as a text (or portion of a text) that stands in a semantic relationship to other texts in terms of the questions it raises and the questions it answers. Muresan (2008) has proposed a conceptualization of meaning where “understanding” a text is the ability to correctly answer, at the conceptual level, all the questions asked about that text. Formally, *Meaning = Text + all Questions/Answers* w.r.t that *Text*. Unlike meaning as truth conditions, where the problem of meaning equivalence is reduced to logical form equivalence, meaning equivalence is reduced to semantic equivalence of DAGs/subDAGs which encode underlying meaning.

The idea that texts both imply questions and provide answers is a way to capture the explicit and implicit content of a text while also indicating actual or potential ties between segments of text within a contribution or between contributions. This conceptualization of meaning for the purposes of computational analysis is compatible with the essential character of argumentative discourse. Meaning-level representation overcomes key limitations in shallow linguistic processing, abstracting away from the surface form of text. The graphical representation treats sentences as vertices and the relation between sentences as edges thus providing a scaffolding for deeper representations of text that make argumentative relations explicit. In addition to labeling types of argumentative relationships, this representation could also identify the underlying perspectives at stake around an issue or the paths that lines of disagreement have taken. The engine for such a ground up representation of discourse from text lies in seeing how segments of text are connected by the kinds of questions and answers they provide, project, and assume. Until now, however, showing argumentative relations graphically has been more in the province of argument visualization techniques.

## 3. VISUALIZATION OF ARGUMENT RELATIONS

Computer supported argument visualization systems (CSAV) provide methods for representing an exchange of messages in terms of the argumentative relations among contributions or portions thereof. CSAVs thus aid interacting collectives in structuring and understanding their collective interaction and reasoning over time. These systems, though quite informative, have yet to leverage the potential for computational systems to detect underlying perspectives, issue formation, and reasoning that happens within an interacting collective. In addition, CSAV systems are built around specific ontologies for representing argumentative relations that, while meant to be general, often serve very particular purposes in articulating argumentation. The preoccupation with designing specific ontologies diverts attention from development of methods for discovering how collectives manage disagreement, and thus, the actual working ontologies interacting collectives employ in targeting and calling-out what is arguable within some domain or activity.

### *3.1 Approaches to argument visualization and their ontologies*

Systems for the computer supported visualization of argument (CSAV) have developed around the use of computing to support the human classification of text in terms of its argumentative purpose (Kirschner, Buckingham-Shum, & Carr, 2003). CSAV systems aim to adequately represent the exchange of contributions in a manner that makes explicit the argumentative relationship among contributions to some ongoing discussion. CSAVs implement various schemes for actors to classify differences of opinion and map lines of disagreement relevant to some decision or matter of discussion. While a variety of actual applications exist, there are at least three approaches to designing ontologies of argumentative relations for the reconstruction of argumentative discourse in support of discussion and decision support.

Visualizing premise-conclusion relations of arguments is one approach. CSAVs such as Rationale, ArguMed, and Carneades are inspired by Toulmin-style argument descriptions for articulating claim-data-warrant as well as rebuttals and refutation relations in what has been contributed to a discussion. Related systems emphasize Walton's (1999) method of "critical questions" for representing argumentative relations. Aruacuria was built to annotate text in a manner that visually represents these relations as critical question of various argument schemes (Chesñevar, McGinnis, Modgil, Rahwan, Reed, Simari, & South, 2006).

Visualizing issue relations emphasizes diagramming how what has been said stands as ideas that answer questions (i.e., issues) for which there are arguments for and against how well the idea answers the question. Such systems reflect Kunz & Rittel's (1970) conceptualization of issue-based information systems that support the articulation of argumentation for the purposes of making sense of complex wicked problems. This approach captures and represents message exchange in an issues-answers-arguments format that articulates rationales for choices. The representation can be consulted during the decision-making process and used later as a historical record. Compendium is a computerized version of this method that enables users to label their contributions as issues, ideas, and arguments. The system renders the annotation as a visual map of the discussion for all participants to see the unfolding contributions as a network representing the collective rationale for the choice. Cohere, Debategraph, and Deliberatorium are web-based systems that provide issue relations annotation for large groups and communities of users.

Visualizing role relations emphasizes the roles actors take up relative to each other in pursuing their differences of opinion around an issue. The most common role relation modeled is the pro-contra relationship between extended contributions made to an ongoing debate. Applications such as Debatepedia and Debate.org provide structured ways for users to make contributions to a defined issue and for a community of participants to develop the argumentation around an issue. The CSAV makes conflicting points of view apparent and reveals lines of disagreement. The coordination is maintained by providing differing roles to contributors, moderators, curators and overhearing audience (e.g., Debatepedia) or to self-manage the development of the debate through tagging likes and replies (e.g. Debate.org).

What these approaches to visualizing argumentative relations reveal is the potential for computing to be used in representing the often complex expressions of differences of opinion and lines of disagreement in the discourse among multiple actors. These approaches attempt to recognize the sequential playing out of argument while articulating the implicit issue structure, presumptions and assumptions, and other non-propositional elements of discourse, (e.g., roles, turns, sequences) that can be important to the management of disagreement.

Common to all of these approaches is the on-the-fly annotation of contributions or text performed by the users of the system. The visualization approaches have been largely dependent on manual coding of argumentative relations. The user either has to make a choice about how to label a contribution to some ongoing discussion while making the contribution or else, post-hoc, a user annotates and labels a portion of text as a particular kind of move relative to something that has been stated. These systems depend on the willingness and ability of the user to apply the labels appropriately. Thus, while offering richer conceptual schemes for modeling argumentative relations between text segments, the scalability of CSAV approaches is limited in the resources it offers for mapping differences of opinion and disagreement expressed in large volumes of text, for decisions involving many actors, or for situations where there is no community to do the coding. Underlying the scalability issue is the fact that, despite the rich ontologies, in practice, the application of these ontologies to text is based largely on ad hoc connections between the annotation, the text, and the implicit issue network. How much faith any individual or collective can have in the relation between the argumentative map and the argumentative terrain is anybody's guess (e.g., Hua & Kimbrough, 1998).

*3.2 Design requirement: Discovering ontologies of argumentative relations*

Argument, from a pragmatic perspective, is both universal and particular in that it is a method for managing differences that are tailored to the substantive problems of human activities and the characteristics of the natural and institutional world in which collectives find themselves. This is an important upshot of Jacobs and Jackson's (1980; 1989) pragmatic argument theory of the local management of disagreement, which can be extended to suggest that ontologies of common sense reasoning could be discovered in an interacting collective's practices of disagreement management. Indeed, Goodwin and Wenzel's (1979) analysis of proverbs about argument suggests such an approach as does Toulmin's (2003) method of discovering fields of argument by attending to which modalities of reasoning are field independent and which are field invariant. Detecting and articulating practices of calling-out would offer genuine insight into the practical reasoning of collectives and also the prospect of creative intervention as communities make sense of circumstances and determine prudent courses of action. The working ontologies for argumentative relations particular to any collective are built from expectations about what counts as premise-conclusion relations, relevance of a contribution to an issue, pragmatic commitments of obligations and rights, and the relevance of a contribution to an activity (Aakhus, 2013). These are the resources from which questions about prior contributions are generated and thus are also a resource for how meaning is worked out. While many questions could be posed about a contribution, not all are posed and thus preferences for patterns of calling-out emerge in different collectives.

In terms of advancing the augmentation of interaction and reasoning, there lurks a more subtle and deep issue about extant CSAV approaches. CSAVs suggest that actors can build taxonomies, or even folksonomies, of complex decision situations, thus helping parties make sense of collective reasoning across complex interactions. The different ontologies that visualization systems rely on for representing differences and disagreement render argumentative representations of the discourse in ways that no doubt vary in their usefulness for augmenting interaction and reasoning. Beyond that, the labels provided by the system for describing behavior, carry their own normative commitments about the purpose of dialogue and the effective and appropriate moves to be made in pursuing differences of opinion and managing disagreement. Even very simple, seemingly unobtrusive systems call for such choices. For instance, debate.org and debatepedia use a pro-con set up that treats issues as having exactly two sides. Indeed, all kinds of ICTs can be understood, in light of the discussion above, as incorporating presuppositions about the conduct of argumentation (e.g., Aakhus, 2002).

Current CSAV approaches are focused more on creating ontologies for certain kinds of argument than on developing methods to discover the ontologies of argument operating within an interacting collective. Discovering how collectives reason involves identifying the way that matters are contested, issues are constructed, and opposition is pursued -- that is, the common practices within a collective for targeting and calling-out aspects of discourse to construct issues and lines of disagreement. The point here is not to disparage CSAVs for their choices in ontology design but to point out the importance of ontology design in developing support systems. Clearly there is value to the development of ontologies that highlight specific aspects of discourse in order to achieve a particular purpose. However this strength of designed ontologies can also be a limiting factor in supporting reasoning in interacting collectives as the implementations transform discourse around the particular normative and descriptive commitments of the ontology. A given ontology can only recognize what it is designed to recognize. Given that argument varies across collectives and the activities in which they are engaged, there is a need to discover the practical ontologies of argumentative relations that collectives work with in making sense of circumstances and determining prudent courses of action. Indeed, the way circumstances are framed and the way problems, choices, and solutions are conceived depends on the tacit network of presumptions and assumptions for any discourse.

## 4. CONCLUSION: DESIGNING RESPONSE CENTERED ARGUMENTATION SUPPORT

Two research endeavors illustrate the potential of combining elements of visualization and NLP in a single system. But even when NLP and visualization systems are ostensibly combined, it is apparent that the two approaches do not yet take into account the phenomenon of disagreement space nor the potential for discovering the practical ontology of interacting collectives.

### *4.1 Illustration*

Pollatta and Delmonte (2011) present a system for analyzing and visually representing interactions via conversation graphs that represent meetings in which decisions are made. The graphs illustrate zones where decision-making took place and summarize who was involved for how many turns. Their system recognizes pairing of speech acts that stand in an argumentative relations to each other (e.g., Propose-Request/Accept/Reject) on the basis of textual cues and shows how a subsequent contribution replies to a prior contribution. They map relation labels such as statement, cause and motivation to five argumentative labels taken from their meeting description schema (MDS) (e.g., Accept, Reject/Disagree).

Pallotta and Delmonte's approach illustrates how NLP and Visualization can be brought together but stops short of illuminating argumentative content and argumentative relations in a broader way. Their graphs do little to visualize specific differences of opinion or lines of disagreement, let alone identify the underlying perspective, common-places or lines of reasoning that point to the tacit network of presumptions and assumptions operating in the discourse. They use a natural language processing system for deep understanding using Lexical Functional Grammar (LFG) (Bresnan, 2001). Grammar formalisms developed for deep linguistic processing such as LFG are not currently known to be learnable from data (neither theoretically, nor empirically), unlike the Lexicalized Well-Founded Grammar described above. Their use in large scale empirical investigations is limited since the grammars need to be handwritten, usually by a large team of grammar engineers and linguists, and are hard to maintain. Adaptation of these grammars to different domains and genres requires substantial effort.

Murakami, Nichols, Mizuno, Watanabe, Masuda, Goto, Ohki, et al. (2010) present a system that uses NLP techniques such as sentiment analysis to produce elegant "statement maps" that group sentences that agree, contradict each other, or provide evidence in support of another sentence. The sentences may come from the same or different documents. Despite the evident usefulness of the statement maps, the paper focuses only on the safety of vaccines and does not provide evidence that the system can be readily generalized to handle other issues.

Their approach resembles the aims of identifying premise-conclusion relations but it appears their approach confounds aspects of premise-conclusion relations expressed or presumed by statements with issue relations that have to do with the action of one contribution on another. From the perspective of argumentation, premise-conclusion relations would be matters of formal logic or practical reasoning schemes like cause, sign, and generalization.

### *4.2 A response centered approach*

In regard to discovering practical ontologies of argumentative relations, both projects reveal the potential of combining NLP and visualization but both are preoccupied with designing an ontology of argumentative relations rather than discovering the argumentative practice of interacting collectives.

While argument in its simplest form is pure opposition it also involves the classic sense of making arguments and is most interesting when having and making arguments occur together in a manner tailored to the circumstances of the collective and the activities in which they are engaged. The previous discussion highlights the need for a response-centered meta-ontology for discovering argumentative relations among contributions to a discourse. In conclusion, then, we suggest some additional high level design requirements for developing socially intelligent systems to augment human reasoning and interaction that follow from the preceding discussion. Such a system should:

- Operate at the level of collective reasoning as generated and performed by groups, organizations, communities and networks of actors.
- Be based on a pragmatic understanding of argument as statements about the world and commitments about action.
- Provide enough formality to exploit the power of computation while avoiding inflexibility that obscures the actual interaction and reasoning practices of collectives.
- Build representations of argumentation sensitive to the surface level playing out of the argument and to the underlying assumptions and overarching perspectives.
- Develop NLP systems able to use both rich representations needed for modeling argumentation and machine learning techniques for robustness and scalability.

REFERENCES

Aakhus, M. (2002). Modeling reconstruction in groupware technology. In F. H. van Eemeren (Ed.), *Advances in pragma-dialectics* (pp. 121–126). Newport News, VA: Vale Press.

Aakhus, M. (2013). Deliberation digitized: Designing disagreement space through communication-information services. *Journal of Argumentation in Context*, *2*(1), 101–126.

Agarwal, A., Biadsy F., & McKeown K. (2009). Contextual phrase-level polarity analysis using lexical affect scoring and syntactic n-grams. *Proceedings of the 12th Conference on the European Chapter of the Association for Computational Linguistics* (EACL2009), Athens, Greece, 2009

Bresnan, J. (2001). *Lexical-functional syntax*. Malden, Mass.: Blackwell.

Bui, T. H., & Peters, S. (2010). Decision detection using hierarchical graphical models. *Proceedings of the ACL 2010 Conference Short Papers* (pp. 307–312). Association for Computational Linguistics.

Bunt, H. (1999). Context representation for dialogue management. In *Modeling and Using Context* (pp. 77-90). Springer Berlin Heidelberg.

Bunt, H., Alexandersson, J., Carletta, J., Choe, J. W., Chengyu Fang, A., Hasida, K., … Traum, D. (2010). Towards an ISO Standard for Dialogue Act Annotation. *In Seventh conference on International Language Resources and Evaluation (LREC'10)*.

Bush, V. (1945). As we may think. *The Atlantic Monthly, 176,* 101-108

Carenini, G., & Murray, G. (2009). Visual structured summaries of human conversations. *CHI* (pp. 5–8). Boston, MA: ACM.

Chesñevar, C., McGinnis, J., Modgil, S., Rahwan, I., Reed, C., Simari, G., South, M., et al. (2006). Towards an argument interchange format. *The Knowledge Engineering Review*, *21*(4), 293–316.

Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. *In Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 168–177). New York, NY, USA: ACM.

Eemeren, F. H. Van, Grootendorst, R., Jackson, S. & Jacobs, S. (1993). *Reconstructing Argumentative Discourse*. Tuscaloosa, AL: University of Alabama Press.

Engelbart, D. C. (1962). *Augmenting human intellect: A conceptual framework*. Stanford, CA.

Galley, M., McKeown, K., Hirschberg, J., & Shriberg, E. (2004). Identifying agreement and disagreement in conversational speech: Use of bayesian networks to model pragmatic dependencies. *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics* (p. 669). Association for Computational Linguistics.

Goodwin, P. D. & Wenzel, J. W. (1979). Proverbs and Practical Reasoning: A study in Socio-Logic. *Quarterly Journal of Speech*, *65*, 289–302.

Hawes, T., Lin, J. & Resnik, P. (2009). Elements of a computational model for multi-party discourse: The turn-taking behavior of Supreme Court justices. *Journal of the American Society for Information Science and Technology*, *60*(8), 1607–1615.

Hillard, D., Ostendorf, M. & Shriberg, E. (2003). Detection of agreement vs. disagreement in meetings: Training with unlabeled data. *Proc. HLT-NAACL Conference*.

Hsueh, P. P., & Moore, J. D. (2007). Automatic Decision Detection in Meeting Speech. *Machine Learning for Multimodal Interaction*, 168–179.

Hua, G. H., & Kimbrough, S. O. (1998). On hypermedia-based argumentation decision support systems. *Decision Support Systems*, *22*(3), 259–275.

Jackson, S., & Jacobs, S. (1980). Structure of conversational argument: Pragmatic bases for the enthymeme. *Quarterly Journal of Speech*, *66*(3), 251–265.

Jacobs, S., & Jackson, S. (1989). Building a model of conversational argument. *Rethinking communication* (Vol. 2, pp. 153–171).

Kim, S. M., & Hovy, E. (2006). Automatic identification of pro and con reasons in online reviews. *Proceedings of the COLING/ACL on Main conference poster sessions* (pp. 483–490). Association for Computational Linguistics.

Kunz, W., & Rittel, H. (1970). *Issues as Elements of lnformation Systems*. Working Paper 131, Institute of Urban and Regional Development, University California, Berkeley.

Licklider, J. C. R. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, *HFE-1*, 4-11.

Murakami, K., Nichols, E., Mizuno, J., Watanabe, Y., Masuda, S., Goto, H., Ohki, M., et al. (2010). Statement Map : Reducing Web Information Credibility Noise through Opinion Classification. *23rd International Conference on Computational Linguistics* (pp. 21-28). Toronto, ON.

Muresan, S & O. Rambow (2007). Grammar Approximation by Representative Sublanguage: A New Model for Language Learning. *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics (ACL 2007)* (pp. 832–839). Prague, Czech Republic.

Muresan, S. (2008). Learning to Map Text to Graph-based Meaning Representations via Grammar Induction. *Proceedings of the 3rd TextGraphs Workshop on Graph-based Algorithms for Natural Language Processing at COLING 2008* (pp. 9-16). Association for Computational Linguistics.

Muresan, S. (2011). Learning for Deep Language Understanding. *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI'11)*, (pp, 1858-1865). AAAI Press.

Muresan, S. (2013). Ontology-based Semantic Interpretation via Grammar Constraints. In Oltramari, A.,Vossen, P., Qin, L., Hovy, E. (Eds.), *New Trends of Research in Ontologies and Lexical Resources (*pp 187- 207). Springer-Verlag.

Pallotta, V., & Delmonte, R. (2011). Automatic argumentative analysis for interaction mining. *Argument & Computation*, *2*(2-3), 77–106.

Pallotta, V., Niekrasz, J., & Purver, M. (2005). Collaborative and argumentative models of meeting discussions. *In Proceeding of CMNA-05 international workshop on Computational Models of Natural Arguments* (part of IJCAI 2005).

Pallotta, V., Seretan, V., & Ailomaa, M. (2007). User requirements analysis for meeting information retrieval based on query elicitation. *Annual Meeting-Association for Computational Linguistics* (Vol. 45, p. 1008).

Pang, B., Lee, L., & Vaithyanathan, S. (2002). Thumbs up?: sentiment classification using machine learning techniques. *In Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing* (EMNLP) (Vol. 10, pp. 79–86). Association for Computational Linguistics.

Popescu, A.-M., & Etzioni, O. (2005). Extracting product features and opinions from reviews. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing* (pp. 339–346).

Sebastiani, F. (2002). Machine learning in automated text categorization. *ACM Computing Survey*, *34*(1), 1–47.

Somasundaran, S., Namata, G., & Getoor, L. (2009). Opinion Graphs for Polarity and Discourse Classification. *Language*, (August), 66–74.

Thomas, M., Pang, B. & Lee, L. (2006). Get out the vote: Determining support or opposition from Congressional floor-debate transcripts. *Proceedings of EMNLP* (pp. 327–335).

Toulmin, S. E. (2003). The uses of argument. Cambridge University Press.

Walton, D. (1999). The New Dialectic: A Method of Evaluating an Argument Used for Some Purpose in a Given Case. *Argument*, *13*, 70–91.

Wilson, T., Wiebe, J., & Hwa, R. (2004). Just how mad are you? Finding strong and weak opinion clauses. In *Proceedings of the National Conference on Artificial Intelligence* (pp. 761–769).

Wilson, T., J. Wiebe, & P. Hoffmann (2005). Recognizing contextual polarity in phrase-level sentiment analysis. *Proceedings of the Joint Human Language Technology Conference and the Conference on Empirical Methods in Natural Language Processing (HLTEMNLP-2005).*