

What's Hot in Human Language Technology: Highlights from NAACL HLT 2015

Joyce Y. Chai

Computer Science and Engineering
Michigan State University
East Lansing, MI 48824, USA
jchai@cse.msu.edu

Anoop Sarkar

Computer Science
Simon Fraser University
Burnaby, BC V5A 1S6, Canada
anoop@sfu.ca

Rada Mihalcea

Computer Science and Engineering
University of Michigan
Ann Arbor, MI 48109, USA
mihalcea@umich.edu

Introduction

The Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technology (NAACL HLT) is a premier conference reporting outstanding research on human language technology. NAACL HLT 2015 set a new record with over seven hundred submissions and close to one thousand participants. The program covered a broad range of topic areas organized into eighteen tracks (Mihalcea, Chai, and Sarkar 2015) that included different levels of language processing (e.g., phonology, morphology, syntax, semantics, discourse and pragmatics) and their applications in information extraction and question answering, information retrieval, generation and summarization, machine translation, sentiment analysis, text categorization and topic modeling, and dialog and interactive systems. It also had special tracks devoted to language resources and evaluations, spoken language processing, linguistic and psycholinguistic aspects of Natural Language Processing (NLP), and machine learning for NLP.

Emerging Research Areas

In this section we give a brief overview on two emerging research areas at NAACL HLT 2015.

Language and Vision

A new track on language and vision was introduced for the first time at NAACL HLT 2015 with an intent to broaden NLP research that is situated in a rich visual and perceptual context. This topic area has received significant attention in our community in the past few years. The keynote talk by Prof. Fei-Fei Li from Stanford University highlighted the importance of language in the quest for visual intelligence and motivated interdisciplinary research in this area. Most contributions in this track centered around the following two research problems:

Alignment between Language Instructions and Video Segments. This task involves aligning natural language descriptions to the corresponding video segments. An approach based on HMM, together with a deep convolutional neural network (CNN) classifier for food detection, was developed and evaluated in a cooking domain (Malmaud et al.

2015). Several discriminative models with latent variables were also explored to learn better alignment models in a wet-lab experiment domain (Naim et al. 2015). As alignment is often the first step in many problems involving language and vision, these approaches and empirical results provide important baselines for future work on this topic.

Video/Image to Text Generation. Watching a video and describing what happens in the video may be an easy task for humans, but it is extremely challenging for computer systems. To help focus on language grounding without being distracted by the noise from object recognition in real images, Ortiz and co-authors applied abstract scenes for text generation (Gilberto Mateos Ortiz, Wolff, and Lapata 2015). The results provide better understanding of the challenges and requirements for text generation from real images. Venugopalan and colleagues presented an approach based on deep recurrent neural networks to generate language descriptions from open-domain videos with large vocabularies (Venugopalan et al. 2015). They have incorporated models learned from a large scale image data to the video domain and resulted in state-of-the-art performance on a benchmark dataset for video to text generation.

NLP for Web, Social Media and Social Sciences

NLP has had an increasingly relevant role in the humanities and social sciences. On the one hand, models of language processing can be used to uncover social dynamics; and on the other hand, large amounts of data available on the web and social media can be used to discover characteristics of language use. The keynote talk by Prof. Lillian Lee from Cornell University highlighted the effect of language use and showed NLP can lead to interesting insights into topics in sociolinguistics and the social sciences such as what makes a memorable phrase, or how different sides of a debate frame the issue through language use. Framing was also explored in several papers in the main program (Baumer et al. 2015; Guerini, Özbal, and Strapparava 2015).

Krishnan and Eisenstein (Krishnan and Eisenstein 2015) presented an unsupervised approach to learn signed social networks, in which edges between individuals are annotated with information about the nature of the relationship by exploiting language, in particular by looking at address terms such as *dude* or *Mr*. This work provided a unification of un-

supervised learning models from NLP with social network analysis models from the social sciences and received one of the two *best student paper* awards at the conference.

Social media provides massive amounts of data, and also provides opportunities for new NLP tasks, e.g. inferring political orientation from language use (Mohammady Ardehaly and Culotta 2015), or learning response generation from Twitter conversations (Sordoni et al. 2015).

Trends in Methodology

Distributional Semantic Models

Distributional approaches model semantics of linguistic items based on the Distributional Hypothesis that words that are used in the same context tend to have similar meanings. These approaches have received increasing attention in recent years as they provide a flexible and scalable means to model inductive learning for word meanings. Several novel contributions on distributional models have been presented at the conference, for example, examining semantic composition (Fyshe et al. 2015), addressing verbs with multiple senses (Greenberg, Sayeed, and Demberg 2015), enriching word vectors with visual features (Lazaridou, Pham, and Baroni 2015), and using substitutes to represent target words (Melamud, Dagan, and Goldberger 2015).

Faruqui et al. developed a new method for incorporating information from semantic resources (e.g., WordNet) into distributional models (Faruqui et al. 2015). Instead of using a joint training objective to incorporate taxonomic semantic knowledge and train distributional word vectors, this work proposes “retrofitting” that adapts the trained vectors post hoc. The results have shown the method is more effective than joint training across a number of common tasks. Since retrofitting can be applied to any word vectors, this work is very promising and has received one of the two *best student paper* awards at the conference.

Beyond modeling word semantics, Soricut and Och have applied a distributional approach to learn morphological transformation rules (Soricut and Och 2015). It accounts for the semantic relations that exist between words belonging to the same morphological family and utilizes transforms between related word forms in a word embedding space to induce morphological relations. The evaluations of this method have shown significant improvements across six different languages. This work has received the *best paper* award at the conference.

Deep Neural Networks

We have also observed an increasing interest in applying deep neural networks (NNs) to many diverse NLP tasks such as document categorization (Johnson and Zhang 2015), video-to-text generation (Venugopalan et al. 2015), word sense disambiguation (Taghipour and Ng 2015), paraphrase identification (Yin and Schütze 2015; Zhao, Hassan, and Auli 2015), and learning response generation from Twitter conversations (Sordoni et al. 2015), etc. Effective discriminative training of NNs was also a focus, with multi-task learning as one of the approaches (Liu et al. 2015).

Deep neural networks have also been very influential in machine translation (MT). Several papers in this area continued the trend of using NNs and continuous representations in MT to model larger bilingual contexts (Xing et al. 2015) and using NNs for reordering the source to minimize issues with diverse word order between source and target in MT (de Gispert, Iglesias, and Byrne 2015).

NLP Tools

NAACL HLT 2015 has also featured over thirty demos (Gerber, Havasi, and Lacatusu 2015) where researchers demonstrated latest development of integrated systems and NLP tools. In particular, several tools addressing semantic and discourse processing have been introduced.

Abstract Meaning Representation (AMR) is a semantic representation language which intends to assign the same graph representation to capture the semantics from sentences with the same basic meaning. This representation has received increasing attention in the NLP community because of its simplicity and expressive power. Vanderwende and colleagues (Vanderwende, Menezes, and Quirk 2015) have demonstrated a parser that can generate AMR not only for English sentences, but also for sentences in French, German, Spanish and Japanese where AMR annotations are not available. A visualization tool is introduced by Saphra and Lopez (Saphra and Lopez 2015) which will allow developers to better explore AMR annotations for cross-language alignments. Beyond semantic processing, Surdeanu and colleagues demonstrated two open-source discourse parsers for Rhetorical Structure Theory (Surdeanu, Hicks, and Valenzuela-Escarcega 2015) from text.

Conclusion

NLP has grown into one of the most exciting and diverse community of researchers, with an ever increasing number of researchers, many big and small companies working in this field, and a vibrant community of learners eager to get prepared to take on some of the fun and exciting challenges in the field. NAACL HLT 2015 was a testimony to the vibrancy and vitality of the field. We have seen several emerging trends in the field, addressing research at the intersection of language and vision, computational social sciences, and work that take advantage of the recent advances in distributional semantic models and deep learning. Details of the activities at the conference venue including papers grouped by topic are archived on the conference web page: <http://naacl.org/naacl-hlt-2015>.

We are only able to give a few examples here. Many interesting and novel contributions were featured at the conference. We encourage the reader to visit the ACL Anthology (<http://www.aclweb.org/anthology/N15-1>) which contains an open-access repository of the NAACL HLT 2015 program. The ACL Anthology (<http://aclweb.org/anthology/>) contains all of the journals, conferences and workshops in NLP since 1965.

References

- Baumer, E.; Elovic, E.; Qin, Y.; Polletta, F.; and Gay, G. 2015. Testing and comparing computational approaches for identifying the language of framing in political news. In *Proceedings of NAACL HLT 2015*, 1472–1482. Denver, Colorado: Association for Computational Linguistics.
- de Gispert, A.; Iglesias, G.; and Byrne, B. 2015. Fast and accurate preordering for smt using neural networks. In *Proceedings of NAACL HLT 2015*, 1012–1017. Denver, Colorado: Association for Computational Linguistics.
- Faruqui, M.; Dodge, J.; Jauhar, S. K.; Dyer, C.; Hovy, E.; and Smith, N. A. 2015. Retrofitting word vectors to semantic lexicons. In *Proceedings of NAACL HLT 2015*, 1606–1615. Denver, Colorado: Association for Computational Linguistics.
- Fyshe, A.; Wehbe, L.; Talukdar, P. P.; Murphy, B.; and Mitchell, T. M. 2015. A compositional and interpretable semantic space. In *Proceedings of NAACL HLT 2015*, 32–41. Denver, Colorado: Association for Computational Linguistics.
- Gerber, M.; Havasi, C.; and Lacatusu, F., eds. 2015. *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*. Denver, Colorado: Association for Computational Linguistics.
- Gilberto Mateos Ortiz, L.; Wolff, C.; and Lapata, M. 2015. Learning to interpret and describe abstract scenes. In *Proceedings of NAACL HLT 2015*, 1505–1515. Denver, Colorado: Association for Computational Linguistics.
- Greenberg, C.; Sayeed, A.; and Demberg, V. 2015. Improving unsupervised vector-space thematic fit evaluation via role-filler prototype clustering. In *Proceedings of NAACL HLT 2015*, 21–31. Denver, Colorado: Association for Computational Linguistics.
- Guerini, M.; Özbal, G.; and Strapparava, C. 2015. Echoes of persuasion: The effect of euphony in persuasive communication. In *Proceedings of NAACL HLT 2015*, 1483–1493. Denver, Colorado: Association for Computational Linguistics.
- Johnson, R., and Zhang, T. 2015. Effective use of word order for text categorization with convolutional neural networks. In *Proceedings of NAACL HLT 2015*, 103–112. Denver, Colorado: Association for Computational Linguistics.
- Krishnan, V., and Eisenstein, J. 2015. you're mr. lebowski, im the dude: Inducing address term formality in signed social networks. In *Proceedings of NAACL HLT 2015*, 1616–1626. Denver, Colorado: Association for Computational Linguistics.
- Lazaridou, A.; Pham, N. T.; and Baroni, M. 2015. Combining language and vision with a multimodal skip-gram model. In *Proceedings of NAACL HLT 2015*, 153–163. Denver, Colorado: Association for Computational Linguistics.
- Liu, X.; Gao, J.; He, X.; Deng, L.; Duh, K.; and Wang, Y.-Y. 2015. Representation learning using multi-task deep neural networks for semantic classification and information retrieval. In *Proceedings of NAACL HLT 2015*, 912–921. Denver, Colorado: Association for Computational Linguistics.
- Malmaud, J.; Huang, J.; Rathod, V.; Johnston, N.; Rabinovich, A.; and Murphy, K. 2015. “what’s cooking?” interpreting cooking videos using text, speech and vision. In *Proceedings of NAACL HLT 2015*, 143–152. Denver, Colorado: Association for Computational Linguistics.
- Melamud, O.; Dagan, I.; and Goldberger, J. 2015. Modeling word meaning in context with substitute vectors. In *Proceedings of NAACL HLT 2015*, 472–482. Denver, Colorado: Association for Computational Linguistics.
- Mihalcea, R.; Chai, J.; and Sarkar, A., eds. 2015. *Proceedings of NAACL HLT 2015*. Denver, Colorado: Association for Computational Linguistics.
- Mohammady Ardehaly, E., and Culotta, A. 2015. Inferring latent attributes of twitter users with label regularization. In *Proceedings of NAACL HLT 2015*, 185–195. Denver, Colorado: Association for Computational Linguistics.
- Naim, I.; Song, Y. C.; Liu, Q.; Huang, L.; Kautz, H.; Luo, J.; and Gildea, D. 2015. Discriminative unsupervised alignment of natural language instructions with corresponding video segments. In *Proceedings of NAACL HLT 2015*, 164–174. Denver, Colorado: Association for Computational Linguistics.
- Saphra, N., and Lopez, A. 2015. Amrica: an amr inspector for cross-language alignments. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, 36–40. Denver, Colorado: Association for Computational Linguistics.
- Sordoni, A.; Galley, M.; Auli, M.; Brockett, C.; Ji, Y.; Mitchell, M.; Nie, J.-Y.; Gao, J.; and Dolan, B. 2015. A neural network approach to context-sensitive generation of conversational responses. In *Proceedings of NAACL HLT 2015*, 196–205. Denver, Colorado: Association for Computational Linguistics.
- Soricut, R., and Och, F. 2015. Unsupervised morphology induction using word embeddings. In *Proceedings of NAACL HLT 2015*, 1627–1637. Denver, Colorado: Association for Computational Linguistics.
- Surdeanu, M.; Hicks, T.; and Valenzuela-Escarcega, M. A. 2015. Two practical rhetorical structure theory parsers. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, 1–5. Denver, Colorado: Association for Computational Linguistics.
- Taghipour, K., and Ng, H. T. 2015. Semi-supervised word sense disambiguation using word embeddings in general and specific domains. In *Proceedings of NAACL HLT 2015*, 314–323. Denver, Colorado: Association for Computational Linguistics.
- Vanderwende, L.; Menezes, A.; and Quirk, C. 2015. An amr parser for english, french, german, spanish and japanese and a new amr-annotated corpus. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, 26–30. Denver, Colorado: Association for Computational Linguistics.
- Venugopalan, S.; Xu, H.; Donahue, J.; Rohrbach, M.; Mooney, R.; and Saenko, K. 2015. Translating videos to natural language using deep recurrent neural networks. In *Proceedings of NAACL HLT 2015*, 1494–1504. Denver, Colorado: Association for Computational Linguistics.
- Xing, C.; Wang, D.; Liu, C.; and Lin, Y. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proceedings of NAACL HLT 2015*, 1006–1011. Denver, Colorado: Association for Computational Linguistics.
- Yin, W., and Schütze, H. 2015. Convolutional neural network for paraphrase identification. In *Proceedings of NAACL HLT 2015*, 901–911. Denver, Colorado: Association for Computational Linguistics.
- Zhao, K.; Hassan, H.; and Auli, M. 2015. Learning translation models from monolingual continuous representations. In *Proceedings of NAACL HLT 2015*, 1527–1536. Denver, Colorado: Association for Computational Linguistics.