CoNLL 2015

The 19th Conference on Computational Natural Language Learning

Proceedings of the Conference

July 30-31, 2015
Beijing, China

Production and Manufacturing by
Taberg Media Group AB
Box 94, 562 02 Taberg
Sweden

CoNLL 2015 Best Paper Sponsors:

ISBN 978-1-941643-77-8

Introduction

The 2015 Conference on Computational Natural Language Learning is the nineteenth in the series of annual meetings organized by SIGNLL, the ACL special interest group on natural language learning. CONLL 2015 will be held in Beijing, China on July 30-31, 2015, in conjunction with ACL-IJCNLP 2015.

For the first time this year, CoNLL 2015 has accepted both long (9 pages of content plus 2 additional pages of references) and short papers (5 pages of content plus 2 additional pages of references). We received 144 submissions in total, of which 81 were long and 61 were short papers, and 17 were eventually withdrawn. Of the remaining 127 papers, 29 long and 9 short papers were selected to appear in the conference program, resulting in an overall acceptance rate of almost 30%. All accepted papers appear here in the proceedings.

As in previous years, CoNLL 2015 has a shared task, this year on Shallow Discourse Parsing. Papers accepted for the shared task are collected in a companion volume of CoNLL 2015.

To fit the paper presentations in a 2-day program, 16 long papers were selected for oral presentation and the remaining 13 long and the 9 short papers were presented as posters. The papers selected for oral presentation are distributed in four main sessions, each consisting of 4 talks. Each of these sessions also includes 3 or 4 spotlights of the long papers selected for the poster session. In contrast, the spotlights for short papers are presented in a single session of 30 minutes. The remaining sessions were used for presenting a selection of 4 shared task papers, two invited keynote speeches and a single poster session, including long, short and shared task papers.

We would like to thank all the authors who submitted their work to CoNLL 2015, as well as the program committee for helping us select the best papers out of many high-quality submissions. We are also grateful to our invited speakers, Paul Smolensky and Eric Xing, who graciously agreed to give talks at CoNLL.

Special thanks are due to the SIGNLL board members, Xavier Carreras and Julia Hockenmaier, for their valuable advice and assistance in putting together this year's program, and to Ben Verhoeven, for redesigning and maintaining the CoNLL 2015 web page. We are grateful to the ACL organization for helping us with the program, proceedings and logistics. Finally, our gratitude goes to our sponsors, Google Inc. and Microsoft Research, for supporting the best paper award and student scholarships at CoNLL 2015.

We hope you enjoy the conference!

Afra Alishahi and Alessandro Moschitti

CoNLL 2015 conference co-chairs

Organizers:

Afra Alishahi, Tilburg University
Alessandro Moschitti, Qatar Computing Research Institute

Invited Speakers:

Paul Smolensky, Johns Hopkins University
Eric Xing, Carnegie Mellon University

Program Committee:

Steven Abney, University of Michigan
Eneko Agirre, University of the Basque Country (UPV/EHU)
Bharat Ram Ambati, University of Edinburgh
Yoav Artzi, University of Washington
Marco Baroni, University of Trento
Alberto Barrón-Cedeño, Qatar Computing Research Institute
Roberto Basili, University of Roma, Tor Vergata
Ulf Brefeld, TU Darmstadt
Claire Cardie, Cornell University
Xavier Carreras, Xerox Research Centre Europe
Asli Celikyilmaz, Microsoft
Daniel Cer, Google
Kai-Wei Chang, University of Illinois
Colin Cherry, NRC
Grzegorz Chrupała, Tilburg University
Alexander Clark, King's College London
Shay B. Cohen, University of Edinburgh
Paul Cook, University of New Brunswick
Bonaventura Coppola, Technical University of Darmstadt
Danilo Croce, University of Roma, Tor Vergata
Aron Culotta, Illinois Institute of Technology
Scott Cyphers, Massachusetts Institute of Technology
Giovanni Da San Martino, Qatar Computing Research Institute
Walter Daelemans, University of Antwerp, CLiPS
Ido Dagan, Bar-Ilan University
Dipanjan Das, Google Inc.
Doug Downey, Northwestern University
Mark Dras, Macquarie University
Jacob Eisenstein, Georgia Institute of Technology
James Fan, IBM Research
Raquel Fernandez, ILLC, University of Amsterdam
Simone Filice, University of Roma "Tor Vergata"
Antske Fokkens, VU Amsterdam
Bob Frank, Yale University
Stefan L. Frank, Radboud University Nijmegen
Dayne Freitag, SRI International
Michael Gamon, Microsoft Research

Wei Gao, Qatar Computing Research Institute
Andrea Gesmundo, Google Inc.
Kevin Gimpel, Toyota Technological Institute at Chicago
Francisco Guzmán, Qatar Computing Research Institute
Thomas Gärtner, University of Bonn and Fraunhofer IAIS
Iryna Haponchyk, University of Trento
James Henderson, Xerox Research Centre Europe
Julia Hockenmaier, University of Illinois Urbana-Champaign
Dirk Hovy, Center for Language Technology, University of Copenhagen
Rebecca Hwa, University of Pittsburgh
Richard Johansson, University of Gothenburg
Shafiq Joty, Qatar Computing Research Institute
Lukasz Kaiser, CNRS & Google
Mike Kestemont, University of Antwerp
Philipp Koehn, Johns Hopkins University
Anna Korhonen, University of Cambridge
Annie Louis, University of Edinburgh
André F. T. Martins, Priberam, Instituto de Telecomunicacoes
Yuji Matsumoto, Nara Institute of Science and Technology
Stewart McCauley, Cornell University
David McClosky, IBM Research
Ryan McDonald, Google
Margaret Mitchell, Microsoft Research
Roser Morante, Vrije Universiteit Amsterdam
Lluís Màrquez, Qatar Computing Research Institute
Preslav Nakov, Qatar Computing Research Institute
Aida Nematzadeh, University of Toronto
Ani Nenkova, University of Pennsylvania
Vincent Ng, University of Texas at Dallas
Joakim Nivre, Uppsala University
Naoaki Okazaki, Tohoku University
Miles Osborne, Bloomberg
Alexis Palmer, Heidelberg University
Andrea Passerini, University of Trento
Daniele Pighin, Google Inc.
Barbara Plank, University of Copenhagen
Thierry Poibeau, LATTICE-CNRS
Sameer Pradhan, cemantix.org
Verónica Pérez-Rosas, University of Michigan
Jonathon Read, School of Computing, Teesside University
Sebastian Riedel, UCL
Alan Ritter, The Ohio State University
Brian Roark, Google Inc.
Dan Roth, University of Illinois
Josef Ruppenhofer, Hildesheim University
Rajhans Samdani, Google Research
Hinrich Schütze, University of Munich
Djamé Seddah, Université Paris Sorbonne
Aliaksei Severyn, University of Trento
Avirup Sil, IBM T.J. Watson Research Center
Khalil Sima'an, ILLC, University of Amsterdam

Kevin Small, Amazon
Vivek Srikumar, University of Utah
Mark Steedman, University of Edinburgh
Anders Søgaard, University of Copenhagen
Christoph Teichmann, University of Potsdam
Kristina Toutanova, Microsoft Research
Ioannis Tsochantaridis, Google Inc.
Jun'ichi Tsujii, Aritificial Intelligence Research Centre at AIST
Kateryna Tymoshenko, University of Trento
Olga Uryupina, University of Trento
Antonio Uva, University of Trento
Antal van den Bosch, Radboud University Nijmegen
Menno van Zaanen, Tilburg University
Aline Villavicencio, Institute of Informatics, Federal University of Rio Grande do Sul
Michael Wiegand, Saarland University
Sander Wubben, Tilburg University
Fabio Massimo Zanzotto, University of Rome "Tor Vergata"
Luke Zettlemoyer, University of Washington
Feifei Zhai, The City University of New York
Min Zhang, SooChow University

Table of Contents

Keynote Talks

On Spectral Graphical Models, and a New Look at Latent Variable Modeling in Natural Language Processing
Eric Xing xx

Does the Success of Deep Neural Network Language Processing Mean – Finally! – the End of Theoretical Linguistics?
Paul Smolensky xxii

Long Papers

A Coactive Learning View of Online Structured Prediction in Statistical Machine Translation
Artem Sokolov, Stefan Riezler and Shay B. Cohen 1

A Joint Framework for Coreference Resolution and Mention Head Detection
Haoruo Peng, Kai-Wei Chang and Dan Roth 12

A Supertag-Context Model for Weakly-Supervised CCG Parser Learning
Dan Garrette, Chris Dyer, Jason Baldridge and Noah A. Smith 22

A Synchronous Hyperedge Replacement Grammar based approach for AMR parsing
Xiaochang Peng, Linfeng Song and Daniel Gildea 32

AIDA2: A Hybrid Approach for Token and Sentence Level Dialect Identification in Arabic
Mohamed Al-Badrashiny, Heba Elfardy and Mona Diab 42

An Iterative Similarity based Adaptation Technique for Cross-domain Text Classification
Himanshu Sharad Bhatt, Deepali Semwal and Shourya Roy 52

Analyzing Optimization for Statistical Machine Translation: MERT Learns Verbosity, PRO Learns Length
Francisco Guzmán, Preslav Nakov and Stephan Vogel 62

Annotation Projection-based Representation Learning for Cross-lingual Dependency Parsing
MIN XIAO and Yuhong Guo 73

Big Data Small Data, In Domain Out-of Domain, Known Word Unknown Word: The Impact of Word Representations on Sequence Labelling Tasks
Lizhen Qu, Gabriela Ferraro, Liyuan Zhou, Weiwei Hou, Nathan Schneider and Timothy Baldwin 83

Contrastive Analysis with Predictive Power: Typology Driven Estimation of Grammatical Error Distributions in ESL
Yevgeni Berzak, Roi Reichart and Boris Katz 94

Cross-lingual syntactic variation over age and gender
Anders Johannsen, Dirk Hovy and Anders Søgaard 103

Cross-lingual Transfer for Unsupervised Dependency Parsing Without Parallel Data
Long Duong, Trevor Cohn, Steven Bird and Paul Cook 113

Detecting Semantically Equivalent Questions in Online User Forums
Dasha Bogdanova, Cicero dos Santos, Luciano Barbosa and Bianca Zadrozny 123

Entity Linking Korean Text: An Unsupervised Learning Approach using Semantic Relations
Youngsik Kim and Key-Sun Choi 132

Incremental Recurrent Neural Network Dependency Parser with Search-based Discriminative Training
Majid Yazdani and James Henderson 142

Instance Selection Improves Cross-Lingual Model Training for Fine-Grained Sentiment Analysis
Roman Klinger and Philipp Cimiano 153

Labeled Morphological Segmentation with Semi-Markov Models
Ryan Cotterell, Thomas Müller, Alexander Fraser and Hinrich Schütze 164

Learning to Exploit Structured Resources for Lexical Inference
Vered Shwartz, Omer Levy, Ido Dagan and Jacob Goldberger 175

Linking Entities Across Images and Text
Rebecka Weegar, Kalle Åström and Pierre Nugues 185

Making the Most of Crowdsourced Document Annotations: Confused Supervised LDA
Paul Felt, Eric Ringger, Jordan Boyd-Graber and Kevin Seppi 194

Multichannel Variable-Size Convolution for Sentence Classification
Wenpeng Yin and Hinrich Schütze 204

Opinion Holder and Target Extraction based on the Induction of Verbal Categories
Michael Wiegand and Josef Ruppenhofer 215

Quantity, Contrast, and Convention in Cross-Situated Language Comprehension
Ian Perera and James Allen 226

Recovering Traceability Links in Requirements Documents
Zeheng Li, Mingrui Chen, LiGuo Huang and Vincent Ng 237

Structural and lexical factors in adjective placement in complex noun phrases across Romance languages
Kristina Gulordava and Paola Merlo 247

Symmetric Pattern Based Word Embeddings for Improved Word Similarity Prediction
Roy Schwartz, Roi Reichart and Ari Rappoport 258

Task-Oriented Learning of Word Embeddings for Semantic Relation Classification
Kazuma Hashimoto, Pontus Stenetorp, Makoto Miwa and Yoshimasa Tsuruoka 268

Temporal Information Extraction from Korean Texts
Young-Seob Jeong, Zae Myung Kim, Hyun-Woo Do, Chae-Gyun Lim and Ho-Jin Choi 279

Transition-based Spinal Parsing
Miguel Ballesteros and Xavier Carreras 289

Short Papers

Accurate Cross-lingual Projection between Count-based Word Vectors by Exploiting Translatable Context Pairs
Shonosuke Ishiwatari, Nobuhiro Kaji, Naoki Yoshinaga, Masashi Toyoda and Masaru Kitsuregawa 300

Deep Neural Language Models for Machine Translation
Thang Luong, Michael Kayser and Christopher D. Manning 305

Finding Opinion Manipulation Trolls in News Community Forums
Todor Mihaylov, Georgi Georgiev and Preslav Nakov 310

Do dependency parsing metrics correlate with human judgments?
Barbara Plank, Héctor Martínez Alonso, Željko Agić, Danijela Merkler and Anders Søgaard 315

Feature Selection for Short Text Classification using Wavelet Packet Transform
Anuj Mahajan, sharmistha Jat and Shourya Roy 321

Learning Adjective Meanings with a Tensor-Based Skip-Gram Model
Jean Maillard and Stephen Clark 327

Model Selection for Type-Supervised Learning with Application to POS Tagging
Kristina Toutanova, Waleed Ammar, Pallavi Choudhury and Hoifung Poon 332

One Million Sense-Tagged Instances for Word Sense Disambiguation and Induction
Kaveh Taghipour and Hwee Tou Ng 338

Reading behavior predicts syntactic categories
Maria Barrett and Anders Søgaard 345

Conference Program

Thursday, July 30, 2015

08:45-09:00 Opening Remarks

09:00-10:10 Session 1.a: Embedding input and output representations

Multichannel Variable-Size Convolution for Sentence Classification
Wenpeng Yin and Hinrich Schütze

Task-Oriented Learning of Word Embeddings for Semantic Relation Classification
Kazuma Hashimoto, Pontus Stenetorp, Makoto Miwa and Yoshimasa Tsuruoka

Symmetric Pattern Based Word Embeddings for Improved Word Similarity Prediction
Roy Schwartz, Roi Reichart and Ari Rappoport

A Coactive Learning View of Online Structured Prediction in Statistical Machine Translation
Artem Sokolov, Stefan Riezler and Shay B. Cohen

10:10-10:30 Session 1.b: Entity Linking (spotlight presentations)

A Joint Framework for Coreference Resolution and Mention Head Detection
Haoruo Peng, Kai-Wei Chang and Dan Roth

Entity Linking Korean Text: An Unsupervised Learning Approach using Semantic Relations
Youngsik Kim and Key-Sun Choi

Linking Entities Across Images and Text
Rebecka Weegar, Kalle Åström and Pierre Nugues

Recovering Traceability Links in Requirements Documents
Zeheng Li, Mingrui Chen, LiGuo Huang and Vincent Ng

10:30-11:00 Coffee Break

11:00-12:00 Session 2.a: Keynote Talk

On Spectral Graphical Models, and a New Look at Latent Variable Modeling in Natural Language Processing
Eric Xing, Carnegie Mellon University

12:00-12:30 Session 2.b: Short Paper Spotlights

Deep Neural Language Models for Machine Translation
Thang Luong, Michael Kayser and Christopher D. Manning

Reading behavior predicts syntactic categories
Maria Barrett and Anders Søgaard

One Million Sense-Tagged Instances for Word Sense Disambiguation and Induction
Kaveh Taghipour and Hwee Tou Ng

Model Selection for Type-Supervised Learning with Application to POS Tagging
Kristina Toutanova, Waleed Ammar, Pallavi Choudhury and Hoifung Poon

Feature Selection for Short Text Classification using Wavelet Packet Transform
Anuj Mahajan, Sharmistha Jat and Shourya Roy

Do dependency parsing metrics correlate with human judgments?
Barbara Plank, Héctor Martínez Alonso, Željko Agić, Danijela Merkler and Anders Søgaard

Learning Adjective Meanings with a Tensor-Based Skip-Gram Model
Jean Maillard and Stephen Clark

Accurate Cross-lingual Projection between Count-based Word Vectors by Exploiting Translatable Context Pairs
Shonosuke Ishiwatari, Nobuhiro Kaji, Naoki Yoshinaga, Masashi Toyoda and Masaru Kitsuregawa

Finding Opinion Manipulation Trolls in News Community Forums
Todor Mihaylov, Georgi Georgiev and Preslav Nakov

12:30-14:00 Lunch Break

14:00-15:30 Session 3: CoNLL Shared Task

The CoNLL-2015 Shared Task on Shallow Discourse Parsing
Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Rashmi Prasad, Christopher Bryant, Attapol Rutherford

A Refined End-to-End Discourse Parser
Jianxiang Wang and Man Lan

The UniTN Discourse Parser in CoNLL 2015 Shared Task: Token-level Sequence Labeling with Argument-specific Models
Evgeny Stepanov, Giuseppe Riccardi and Ali Orkan Bayer

The SoNLP-DP System in the CoNLL-2015 shared Task
Fang Kong, Sheng Li and Guodong Zhou

15:30-16:00 Coffee Break

16:00-17:10 Session 4.a: Syntactic Parsing

A Supertag-Context Model for Weakly-Supervised CCG Parser Learning
Dan Garrette, Chris Dyer, Jason Baldridge and Noah A. Smith

Transition-based Spinal Parsing
Miguel Ballesteros and Xavier Carreras

Cross-lingual Transfer for Unsupervised Dependency Parsing Without Parallel Data
Long Duong, Trevor Cohn, Steven Bird and Paul Cook

Incremental Recurrent Neural Network Dependency Parser with Search-based Discriminative Training
Majid Yazdani and James Henderson

17:10-17:30 Session 4.b: Cross-language studies (spotlight presentations)

Structural and lexical factors in adjective placement in complex noun phrases across Romance languages
Kristina Gulordava and Paola Merlo

Instance Selection Improves Cross-Lingual Model Training for Fine-Grained Sentiment Analysis
Roman Klinger and Philipp Cimiano

Annotation Projection-based Representation Learning for Cross-lingual Dependency Parsing
Min Xiao and Yuhong Guo

Friday, July 31, 2015

09:00-10:10 Session 5.a: Semantics

Detecting Semantically Equivalent Questions in Online User Forums
Dasha Bogdanova, Cicero dos Santos, Luciano Barbosa and Bianca Zadrozny

An Iterative Similarity based Adaptation Technique for Cross-domain Text Classification
Himanshu Sharad Bhatt, Deepali Semwal and Shourya Roy

Making the Most of Crowdsourced Document Annotations: Confused Supervised LDA
Paul Felt, Eric Ringger, Jordan Boyd-Graber and Kevin Seppi

Big Data Small Data, In Domain Out-of Domain, Known Word Unknown Word: The Impact of Word Representations on Sequence Labelling Tasks
Lizhen Qu, Gabriela Ferraro, Liyuan Zhou, Weiwei Hou, Nathan Schneider and Timothy Baldwin

10:10-10:30 Session 5.b: Extraction and Labeling (spotlight presentations)

Labeled Morphological Segmentation with Semi-Markov Models
Ryan Cotterell, Thomas Müller, Alexander Fraser and Hinrich Schütze

Opinion Holder and Target Extraction based on the Induction of Verbal Categories
Michael Wiegand and Josef Ruppenhofer

Temporal Information Extraction from Korean Texts
Young-Seob Jeong, Zae Myung Kim, Hyun-Woo Do, Chae-Gyun Lim and Ho-Jin Choi

10:30-11:00 Coffee Break

11:00-12:00 Session 6.a: Keynote Talk

Does the Success of Deep Neural Network Language Processing Mean – Finally! – the End of Theoretical Linguistics?
Paul Smolensky, Johns Hopkins University

12:00-12:30 Session 6.b: Business Meeting

12:30-14:00 Lunch Break

14:00-15:10 Session 7.a: Language Structure

Cross-lingual syntactic variation over age and gender
Anders Johannsen, Dirk Hovy and Anders Søgaard

A Synchronous Hyperedge Replacement Grammar based approach for AMR parsing
Xiaochang Peng, Linfeng Song and Daniel Gildea

Contrastive Analysis with Predictive Power: Typology Driven Estimation of Grammatical Error Distributions in ESL
Yevgeni Berzak, Roi Reichart and Boris Katz

AIDA2: A Hybrid Approach for Token and Sentence Level Dialect Identification in Arabic
Mohamed Al-Badrashiny, Heba Elfardy and Mona Diab

15:10-15:30 Session 7.b: CoNLL Mix (spotlight presentations)

Analyzing Optimization for Statistical Machine Translation: MERT Learns Verbosity, PRO Learns Length
Francisco Guzmán, Preslav Nakov and Stephan Vogel

Learning to Exploit Structured Resources for Lexical Inference
Vered Shwartz, Omer Levy, Ido Dagan and Jacob Goldberger

Quantity, Contrast, and Convention in Cross-Situated Language Comprehension
Ian Perera and James Allen

15:30-16:00 Coffee Break

16:00-17:30 Session 8.a: Joint Poster Presentation (long, short and shared task papers)

Long Papers

A Joint Framework for Coreference Resolution and Mention Head Detection
Haoruo Peng, Kai-Wei Chang and Dan Roth

Entity Linking Korean Text: An Unsupervised Learning Approach using Semantic Relations
Youngsik Kim and Key-Sun Choi

Linking Entities Across Images and Text
Rebecka Weegar, Kalle Åström and Pierre Nugues

Recovering Traceability Links in Requirements Documents
Zeheng Li, Mingrui Chen, LiGuo Huang and Vincent Ng

Structural and lexical factors in adjective placement in complex noun phrases across Romance languages
Kristina Gulordava and Paola Merlo

Instance Selection Improves Cross-Lingual Model Training for Fine-Grained Sentiment Analysis
Roman Klinger and Philipp Cimiano

Annotation Projection-based Representation Learning for Cross-lingual Dependency Parsing
Min Xiao and Yuhong Guo

Labeled Morphological Segmentation with Semi-Markov Models
Ryan Cotterell, Thomas Müller, Alexander Fraser and Hinrich Schütze

Opinion Holder and Target Extraction based on the Induction of Verbal Categories
Michael Wiegand and Josef Ruppenhofer

Temporal Information Extraction from Korean Texts
Young-Seob Jeong, Zae Myung Kim, Hyun-Woo Do, Chae-Gyun Lim and Ho-Jin Choi

Analyzing Optimization for Statistical Machine Translation: MERT Learns Verbosity, PRO Learns Length
Francisco Guzmán, Preslav Nakov and Stephan Vogel

Learning to Exploit Structured Resources for Lexical Inference
Vered Shwartz, Omer Levy, Ido Dagan and Jacob Goldberger

Quantity, Contrast, and Convention in Cross-Situated Language Comprehension
Ian Perera and James Allen

Short Papers

Deep Neural Language Models for Machine Translation
Thang Luong, Michael Kayser and Christopher D. Manning

Reading behavior predicts syntactic categories
Maria Barrett and Anders Søgaard

One Million Sense-Tagged Instances for Word Sense Disambiguation and Induction
Kaveh Taghipour and Hwee Tou Ng

Model Selection for Type-Supervised Learning with Application to POS Tagging
Kristina Toutanova, Waleed Ammar, Pallavi Choudhury and Hoifung Poon

Feature Selection for Short Text Classification using Wavelet Packet Transform
Anuj Mahajan, Sharmistha Jat and Shourya Roy

Do dependency parsing metrics correlate with human judgments?
Barbara Plank, Héctor Martínez Alonso, Željko Agić, Danijela Merkler and Anders Søgaard

Learning Adjective Meanings with a Tensor-Based Skip-Gram Model
Jean Maillard and Stephen Clark

Accurate Cross-lingual Projection between Count-based Word Vectors by Exploiting Translatable Context Pairs
Shonosuke Ishiwatari, Nobuhiro Kaji, Naoki Yoshinaga, Masashi Toyoda and Masaru Kitsuregawa

Finding Opinion Manipulation Trolls in News Community Forums
Todor Mihaylov, Georgi Georgiev and Preslav Nakov

Shared Task Papers

A Hybrid Discourse Relation Parser in CoNLL 2015
Sobha Lalitha Devi, Sindhuja Gopalan, Lakshmi S, Pattabhi RK Rao, Vijay Sundar Ram and Malarkodi C.S.

A Minimalist Approach to Shallow Discourse Parsing and Implicit Relation Recognition
Christian Chiarcos and Niko Schenk

A Shallow Discourse Parsing System Based On Maximum Entropy Model
Jia Sun, Peijia Li, Weiqun Xu and Yonghong Yan

Hybrid Approach to PDTB-styled Discourse Parsing for CoNLL-2015
Yasuhisa Yoshida, Katsuhiko Hayashi, Tsutomu Hirao and Masaaki Nagata

Improving a Pipeline Architecture for Shallow Discourse Parsing
Yangqiu Song, Haoruo Peng, Parisa Kordjamshidi, Mark Sammons and Dan Roth

JAIST: A two-phase machine learning approach for identifying discourse relations in newswire texts
Son Nguyen, Quoc Ho and Minh Nguyen

Shallow Discourse Parsing Using Constituent Parsing Tree
Changge Chen, Peilu Wang and Hai Zhao

Shallow Discourse Parsing with Syntactic and (a Few) Semantic Features
Shubham Mukherjee, Abhishek Tiwari, Mohit Gupta and Anil Kumar Singh

The CLaC Discourse Parser at CoNLL-2015
Majid Laali, Elnaz Davoodi and Leila Kosseim

The DCU Discourse Parser for Connective, Argument Identification and Explicit Sense Classification
Longyue Wang, Chris Hokamp, Tsuyoshi Okita, Xiaojun Zhang and Qun Liu

The DCU Discourse Parser: A Sense Classification Task
Tsuyoshi Okita, Longyue Wang and Qun Liu

17:30-17:45 Session 8.b: Best Paper Award and Closing

Keynote Talk

On Spectral Graphical Models, and a New Look at Latent Variable Modeling in Natural Language Processing

Eric Xing

Carnegie Mellon University

epxing@cs.cmu.edu

Abstract

Latent variable and latent structure modeling, as widely seen in parsing systems, machine translation systems, topic models, and deep neural networks, represents a key paradigm in Natural Language Processing, where discovering and leveraging syntactic and semantic entities and relationships that are not explicitly annotated in the training set provide a crucial vehicle to obtain various desirable effects such as simplifying the solution space, incorporating domain knowledge, and extracting informative features. However, latent variable models are difficult to train and analyze in that, unlike fully observed models, they suffer from non-identifiability, non-convexity, and over-parameterization, which make them often hard to interpret, and tend to rely on local-search heuristics and heavy manual tuning.

In this talk, I propose to tackle these challenges using spectral graphical models (SGM), which view latent variable models through the lens of linear algebra and tensors. I show how SGMs exploit the connection between latent structure and low rank decomposition, and allow one to develop models and algorithms for a variety of latent variable problems, which unlike traditional techniques, enjoy provable guarantees on correctness and global optimality, can straightforwardly incorporate additional modern techniques such as kernels to achieve more advanced modeling power, and empirically offer a 1-2 orders of magnitude speed up over existing methods while giving comparable or better performance.

This is joint work with Ankur Parikh, Carnegie Mellon University.

Biography of the Speaker

Dr. Eric Xing is a Professor of Machine Learning in the School of Computer Science at Carnegie Mellon University, and Director of the CMU/UPMC Center for Machine Learning and Health. His principal research interests lie in the development of machine learning and statistical methodology, and large-scale computational system and architecture; especially for solving problems involving automated learning, reasoning, and decision-making in high-dimensional, multimodal, and dynamic possible worlds in artificial, biological, and social systems. Professor Xing received a Ph.D. in Molecular Biology from Rutgers University, and another Ph.D. in Computer Science from UC Berkeley. He servers (or served) as an associate editor of the Annals of Applied Statistics (AOAS), the Journal of American Statistical Association (JASA), the IEEE Transaction of Pattern Analysis and Machine Intelligence (PAMI), the PLoS Journal of Computational Biology, and an Action Editor of the Machine Learning Journal (MLJ), the Journal of Machine Learning Research (JMLR). He was a member of the DARPA Information Science and Technology (ISAT) Advisory Group, a recipient of the NSF Career Award, the Sloan Fellowship, the United States Air Force Young Investigator Award, and the IBM Open Collaborative Research Award. He is the Program Chair of ICML 2014.

Keynote Talk

Does the Success of Deep Neural Network Language Processing Mean – Finally! – the End of Theoretical Linguistics?

Paul Smolensky

Johns Hopkins University

smolensky@jhu.edu

Abstract

Statistical methods in natural-language processing that rest on heavily empirically-based language learning – especially those centrally deploying neural networks – have witnessed dramatic improvement in the past few years, and their success restores the urgency of understanding the relationship between (i) these neural/statistical language systems and (ii) the view of linguistic representation, processing, and structure developed over centuries within theoretical linguistics.

Two hypotheses concerning this relationship arise from our own mathematical and experimental results from past work, which we will present. These hypotheses can guide – we will argue – important future research in the seemingly sizable gap separating computational linguistics from linguistic theories of human language acquisition. These hypotheses are:

1. The internal representational format used in deep neural networks for language – numerical vectors – is covertly an implementation of a system of discrete, symbolic, structured representations which are processed so as to optimally meet the demands of a symbolic grammar recognizable from the perspective of theoretical linguistics.
2. It will not be successes but rather the *failures* of future machine learning approaches to language acquisition which will be most telling for determining whether such approaches capture the crucial limitations on human language learning – limitations, documented in recent artificial-grammar-learning experimental results, which support the nativist Chomskian hypothesis asserting that
    - reliably and efficiently learning human grammars from available evidence requires
    - that the hypothesis space entertained by the child concerning the set of possible (or likely) human languages
    - be limited by abstract, structure-based constraints;
    - these constraints can then also explain (in principle at least) the many robustly-respected universals observed in cross-linguistic typology.

This is joint work with Jennifer Culbertson, University of Edinburgh.

Biography of the Speaker

Paul Smolensky is the Krieger-Eisenhower Professor of Cognitive Science at Johns Hopkins University in Baltimore, Maryland, USA. He studies the mutual implications between the theories of neural computation and of universal grammar and has published in distinguished venues including Science and the Proceedings of the National Academy of Science USA. He received the David E. Rumelhart Prize for Outstanding Contributions to the Formal Analysis of Human Cognition (2005), the Chaire de Recherche Blaise Pascal (2008—9), and the Sapir Professorship of the Linguistic Society of America (2015). Primary results include:

- Contradicting widely-held convictions, (i) structured symbolic and (ii) neural network models of cognition are mutually compatible: formal descriptions of the same systems, the mind/brain, at (i) a highly abstract, and (ii) a more physical, level of description. His article "On the proper treatment of connectionism" (1988) was until recently one of the 10-most cited articles in The Behavioral and Brain Sciences, itself the most-cited journal of all the behavioral sciences.
- That the theory of neural computation can in fact strengthen the theory of universal grammar is attested by the revolutionary impact in theoretical linguistics (within phonology in particular) of Optimality Theory, a neural- network-derived symbolic grammar formalism that he developed with Alan Prince (in a book widely released 1993, officially published 2004).
- The learnability theory for Optimality Theory was founded at nearly the same time as the theory itself, in joint work of Smolensky and his PhD student Bruce Tesar (TR 1993; article in the premier linguistic theory journal, Linguistic Inquiry 1998; MIT Press book 2000). This work laid the foundation upon which rests most of the flourishing formal theory of learning in Optimality Theory.
- There is considerable power in formalizing neural network computation as statistical inference/optimization within a dynamical system. Smolensky's Harmony Theory (1981–6) analyzed network computation as Harmony Maximization (an independently-developed homologue to Hopfield's "energy minimization" formulation) and first deployed principles of statistical inference for processing and learning in the bipartite network structure later to be known as the 'Restricted Boltzmann Machine' in the initial work on deep neural network learning (Hinton et al., 2006–).
- Powerful recursive symbolic computation can be achieved with massive parallelism in neural networks designed to process tensor product representations (TR 1987; journal article in Artificial Intelligence 1990). Related uses of the tensor product to structure numerical vectors is currently under rapid development in the field of distributional vector semantics.

Most recently, as argued in Part 1 of the talk, his work shows the value for theoretical and psycholinguistics of representations that share both the discrete structure of symbolic representations and the continuous variation of activity levels in neural network representations (initial results in an article in Cognitive Science by Smolensky, Goldrick & Mathis 2014).