

CoNLL 2018

**The 22nd Conference on
Computational Natural Language Learning**

Proceedings of the Conference

October 31 - November 1, 2018
Brussels, Belgium

Sponsors



©2018 The Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-948087-72-8

Introduction

The 2018 Conference on Computational Natural Language Learning (CoNLL) is the 22nd in the series of annual meetings organized by SIGNLL, the ACL special interest group on natural language learning. CoNLL 2018 will be held on October 31 - November 1, 2018, and is co-located with the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP) in Brussels, Belgium.

CoNLL 2018 followed the tradition of previous CoNLL conferences in inviting only long papers, in order to accommodate papers with experimental material and detailed analysis. The final, camera-ready submissions were allowed a maximum of nine content pages plus unlimited pages of references and supplementary material.

CoNLL 2018 received a record number of 295 submissions in total, out of which 5 had to be rejected for formal reasons and 16 were withdrawn by the authors during the review period. Of the remaining 274 papers, 57 papers were chosen to appear in the conference program, with an overall acceptance rate of 20.8%. One of these was withdrawn after the notification, resulting in 56 papers for the final program: 16 were selected for oral presentation, and the remaining 40 for poster presentation plus lightning oral presentation. All 56 papers appear here in the conference proceedings.

CoNLL 2018 features two invited speakers, Asifa Majid (University of York) and Max Welling (University of Amsterdam / CIFAR). As in recent years, it also features two shared tasks: one on Universal Morphological Reinflection and one on Multilingual Parsing from Raw Text to Universal Dependencies. Papers accepted for the shared tasks are published in companion volumes of the CoNLL 2018 proceedings.

We would like to thank all the authors who submitted their work to CoNLL 2018, and the program committee for helping us select the best papers out of many high-quality submissions. We are grateful to the many program committee members who did a thorough job reviewing our submissions. Due to the growing size of the conference, we also had area chairs, for the first time, supporting the CoNLL organization. We were fortunate to have 12 excellent areas chairs who assisted us greatly in selecting the best programme:

Marine Carpuat, University of Maryland, USA
Paul Cook, University of New Brunswick, USA
Vera Demberg, Saarland University, Germany
Graham Neubig, Carnegie Mellon University, USA
Sebastian Pado, University of Stuttgart, Germany
Siva Reddy, Stanford University, USA
Roi Reichart, Technion, Israel
Alan Ritter, Ohio State University, USA
Tim Rocktäschel, University of Oxford, UK
Mehrnoosh Sadzadeh, Queen Mary University of London, UK
Sameer Singh, University of California, Irvine USA
Yulia Tsvetkov, Carnegie Mellon University, USA

We are immensely thankful to Julia Hockenmaier and to the members of the SIGNLL board for their valuable advice and assistance in putting together this year's program. We also thank Ben Verhoeven, for maintaining the CoNLL 2018 website and Miikka Silfverberg for preparing the proceedings for the main conference. Finally, we would like to thank our hard working assistants, Phong Le and Edoardo Ponti, for their great support with the conference administration and publicity.

Finally, our gratitude goes to our sponsors, Google Inc. and Textkernel, for supporting the conference financially.

We hope you enjoy the conference!

Anna Korhonen and Ivan Titov

CoNLL 2018 conference co-chairs

Conference chairs:

Anna Korhonen, University of Cambridge
Ivan Titov, University of Edinburgh / University of Amsterdam

Invited speakers:

Asifa Maji, University of York
Max Welling, University of Amsterdam / CIFAR

Area chairs:

Marine Carpuat, University of Maryland, USA
Paul Cook, University of New Brunswick, USA
Vera Demberg, Saarland University, Germany
Graham Neubig, Carnegie Mellon University, USA
Sebastian Pado, University of Stuttgart, Germany
Siva Reddy, Stanford University, USA
Roi Reichart, Technion, Israel
Alan Ritter, Ohio State University, USA
Tim Rocktäschel, University of Oxford, UK
Mehrnoosh Sadrzadeh, Queen Mary University of London, UK
Sameer Singh, University of California, Irvine USA
Yulia Tsvetkov, Carnegie Mellon University, USA

Publication chair:

Miikka Silfverberg, University of Colorado, USA

Administration chair:

Phong Le, University of Edinburgh, UK

Publicity / sponsorship chair:

Edoardo Ponti, University of Cambridge, UK

Program Committee:

Omri Abend, Željko Agić, Roei Aharoni, Alan Akbik, Afra Alishahi, Stefanos Angelidis, Marianna Apidianaki, Kenji Araki, Wilker Aziz, Ashutosh Baheti, Fan Bai, Simon Baker, Omid Bakhshandeh, Miguel Ballesteros, Joost Bastings, Timo Baumann, Yonatan Belinkov, Iz Beltagy, Anja Belz, Jonathan Berant, Taylor Berg-Kirkpatrick, Raffaella Bernardi, Yevgeni Berzak, Alexandra Birch, Arianna Bisazza, Yonatan Bisk, Philippe Blache, Bernd Bohnet, Gerlof Bouma, Samuel Bowman, S.R.K. Branavan, Chloé Braud, Daniel Cer, Snigdha Chaturvedi, Boxing Chen, Danqi Chen, Xinchu Chen, Jianpeng Cheng, Christos Christodoulopoulos, Grzegorz Chrupała, Shay B. Cohen, Trevor Cohn, Danish Contractor, Caio Corro, Marta R. Costa-jussà, Ryan Cotterell, Tim Van de Cruys, Joachim Daiber, Amitava Das, Rajarshi Das, Steve DeNeefe, Thomas Demeester, Pascal Denis, Shuoyang Ding, Georgiana Dinu, Simon Dobnik, Jesse Dodge, Li Dong, Doug Downey, Gabriel Doyle, Markus Dreyer, Maximillian Droog-Hayes, Xinya Du, Kevin Duh, Nadir Durrani, Michael Elhadad, Micha Elsner, Federico Fancellu, Afsaneh Fazly, Marcello Federico, Christian Federmann, Yansong Feng, Raquel Fernández, Katja Filippova,

Andrea K. Fischer, Jeffrey Flanigan, Annie Foret, George Foster, Stefan L. Frank, Stella Frank, Dan Garrette, Mehdi Ghanimifard, Dafydd Gibbon, Daniel Gildea, Roxana Girju, Yoav Goldberg, Dan Goldwasser, Matthew R. Gormley, Yvette Graham, Edward Grefenstette, Nitish Gupta, Gholamreza Haffari, Keith Hall, William L. Hamilton, Kazuma Hashimoto, Hua He, Luheng He, Aurélie Herbelot, Dirk Hovy, Renfen Hu, Tim Hunter, Alvin Grissom II, Marco Idiart, Srinivasan Iyer, Laura Jehl, Robin Jia, Anders Johannsen, Alexander Johansen, Richard Johansson, Mandar Joshi, Katharina Kann, Dimitri Kartsaklis, Daisuke Kawahara, Tushar Khot, Yoon Kim, Milton King, Eliyahu Kiperwasser, Svetlana Kiritchenko, Sigrid Klerke, Roman Klinger, Thomas Kober, Philipp Koehn, Ioannis Konstas, Julia Kreutzer, Jayant Krishnamurthy, Germán Kruszewski, Marco Kuhlmann, Sachin Kumar, Jonathan K. Kummerfeld, Oier Lopez de Lacalle, Ni Lao, Yoong Keok Lee, Jochen L. Leidner, Roger Levy, Martha Lewis, Xi Victoria Lin, Xiao Ling, Wei Lu, Suresh Manandhar, Diego Marcheggiani, Yuval Marton, Diana McCarthy, David McClosky, Stephen McGregor, R. Thomas McCoy, Pasquale Minervini, Jeff Mitchell, Daichi Mochihashi, Marie-Francine Moens, Nasrin Mostafazadeh, Smaranda Muresan, Kenton Murray, Ajay Nagesh, Preslav Nakov, Jason Naradowsky, Shashi Narayan, Massimo Nicosia, Xing Niu, Brendan O'Connor, Kemal Oflazer, Ulrike Pado, Alexis Palmer, Martha Palmer, Alexander Panchenko, Ali Hakimi Parizi, Panupong Pasupat, Michael J. Paul, Laura Perez-Beltrachini, Hieu Pham, Yuval Pinter, Barbara Plank, Massimo Poesio, Ella Rabinovich, Preethi Raghavan, Afshin Rahimi, Carlos Ramisch, Marek Rei, Steffen Remus, Xiang Ren, Sebastian Riedel, Martin Riedl, Laura Rimell, Brian Roark, Laurent Romary, Michael Roth, Joseph Le Roux, Alla Rozovskaya, Sebastian Ruder, Andreas Rücklé, Kenji Sagae, Ivan Sanchez, Ryohei Sasano, Carolina Scarton, Christian Scheible, Marten van Schijndel, William Schuler, Philip Schulz, Roy Schwartz, Djamé Seddah, Abigail See, Rico Sennrich, Minjoon Seo, Ehsan Shareghi, Wei Shi, Carina Silberer, Kevin Small, Daniil Sorokin, Vivek Srikumar, Miloš Stanojević, Gabriel Stanovsky, Mark Steedman, Pontus Stenetorp, Karl Stratos, Yu Su, Huan Sun, Partha Talukdar, Christoph Teichmann, Oren Tsur, Oscar Täckström, Patrick Verga, Aline Villavicencio, Andreas Vlachos, Svitlana Volkova, Sabine Schulte im Walde, William Yang Wang, Taro Watanabe, Julie Weeds, Dirk Weissenborn, Johannes Welbl, Michael Wiegand, Gijs Wijnholds, Dekai Wu, Kun Xu, Mohamed Yahya, Bishan Yang, Weiwei Yang, François Yvon, Congle Zhang, Victor Zhong, Shi Zong, Willem Zuidema

Table of Contents

<i>Inductive Bias in Deep Learning</i> Max Welling	xviii
<i>Semantic Spaces Across Diverse Languages</i> Asifa Majid	xix
<i>Embedded-State Latent Conditional Random Fields for Sequence Labeling</i> Dung Thai, Sree Harsha Ramesh, Shikhar Murty, Luke Vilnis and Andrew McCallum	1
<i>Continuous Word Embedding Fusion via Spectral Decomposition</i> Tianfan Fu, Cheng Zhang and Stephan Mandt	11
<i>Dual Latent Variable Model for Low-Resource Natural Language Generation in Dialogue Systems</i> Van-Khanh Tran and Le-Minh Nguyen	21
<i>A Trio Neural Model for Dynamic Entity Relatedness Ranking</i> Tu Nguyen, Tuan Tran and Wolfgang Nejdl	31
<i>A Unified Neural Network Model for Geolocating Twitter Users</i> Mohammad Ebrahimi, Elaheh ShafieiBavani, Raymond Wong and Fang Chen	42
<i>Corpus-Driven Thematic Hierarchy Induction</i> Iliia Kuznetsov and Iryna Gurevych	54
<i>Adversarially Regularising Neural NLI Models to Integrate Logical Background Knowledge</i> Pasquale Minervini and Sebastian Riedel	65
<i>From Strings to Other Things: Linking the Neighborhood and Transposition Effects in Word Reading</i> Stephan Tulkens, Dominiek Sandra and Walter Daelemans	75
<i>Global Attention for Name Tagging</i> Boliang Zhang, Spencer Whitehead, Lifu Huang and Heng Ji	86
<i>Pervasive Attention: 2D Convolutional Neural Networks for Sequence-to-Sequence Prediction</i> Maha Elbayad, Laurent Besacier and Jakob Verbeek	97
<i>Comparing Attention-Based Convolutional and Recurrent Neural Networks: Success and Limitations in Machine Reading Comprehension</i> Matthias Blohm, Glorianna Jagfeld, Ekta Sood, Xiang Yu and Ngoc Thang Vu	108
<i>Uncovering Code-Mixed Challenges: A Framework for Linguistically Driven Question Generation and Neural Based Question Answering</i> Deepak Gupta, Pabitra Lenka, Asif Ekbal and Pushpak Bhattacharyya	119
<i>Learning to Embed Semantic Correspondence for Natural Language Understanding</i> Sangkeun Jung, Jinsik Lee and Jiwon Kim	131
<i>Commonsense Knowledge Base Completion and Generation</i> Itsumi Saito, Kyosuke Nishida, Hisako Asano and Junji Tomita	141
<i>Active Learning for Interactive Neural Machine Translation of Data Streams</i> Álvaro Peris and Francisco Casacuberta	151

<i>Churn Intent Detection in Multilingual Chatbot Conversations and Social Media</i> Christian Abbet, Meryem M’hamdi, Athanasios Giannakopoulos, Robert West, Andreea Hossmann, Michael Baeriswyl and Claudiu Musat	161
<i>Learning Text Representations for 500K Classification Tasks on Named Entity Disambiguation</i> Ander Barrena, Aitor Soroa and Eneko Agirre	171
<i>Hierarchical Attention Based Position-Aware Network for Aspect-Level Sentiment Analysis</i> Lishuang Li, Yang Liu and AnQiao Zhou	181
<i>Bidirectional Generative Adversarial Networks for Neural Machine Translation</i> Zhirui Zhang, Shujie Liu, Mu Li, Ming Zhou and Enhong Chen	190
<i>Latent Entities Extraction: How to Extract Entities that Do Not Appear in the Text?</i> Eylon Shoshan and Kira Radinsky	200
<i>Generalizing Procrustes Analysis for Better Bilingual Dictionary Induction</i> Yova Kementchedjhieva, Sebastian Ruder, Ryan Cotterell and Anders Søgaard	211
<i>Simple Unsupervised Keyphrase Extraction using Sentence Embeddings</i> Kamil Bennani-Smires, Claudiu Musat, Andreea Hossmann, Michael Baeriswyl and Martin Jaggi	221
<i>A Temporally Sensitive Submodularity Framework for Timeline Summarization</i> Sebastian Martschat and Katja Markert	230
<i>Chinese Poetry Generation with a Salient-Clue Mechanism</i> Xiaoyuan Yi, Ruoyu Li and Maosong Sun	241
<i>Multi-Modal Sequence Fusion via Recursive Attention for Emotion Recognition</i> Rory Beard, Ritwik Das, Raymond W. M. Ng, P. G. Keerthana Gopalakrishnan, Luka Eerens, Pawel Swietojanski and Ondrej Miksik	251
<i>Using Sparse Semantic Embeddings Learned from Multimodal Text and Image Data to Model Human Conceptual Knowledge</i> Steven Derby, Paul Miller, Brian Murphy and Barry Devereux	260
<i>Similarity Dependent Chinese Restaurant Process for Cognate Identification in Multilingual Wordlists</i> Taraka Rama	271
<i>Uncovering Divergent Linguistic Information in Word Embeddings with Lessons for Intrinsic and Ex- trinsic Evaluation</i> Mikel Artetxe, Gorka Labaka, Inigo Lopez-Gazpio and Eneko Agirre	282
<i>Comparing Models of Associative Meaning: An Empirical Investigation of Reference in Simple Lan- guage Games</i> Judy Hanwen Shen, Matthias Hofer, Bjarke Felbo and Roger Levy	292
<i>Sequence Classification with Human Attention</i> Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei and Anders Søgaard	302
<i>Sentence-Level Fluency Evaluation: References Help, But Can Be Spared!</i> Katharina Kann, Sascha Rothe and Katja Filippova	313
<i>Predefined Sparseness in Recurrent Sequence Models</i> Thomas Demeester, Johannes Deleu, Frédéric Godin and Chris Develder	324

<i>Learning to Actively Learn Neural Machine Translation</i>	
Ming Liu, Wray Buntine and Gholamreza Haffari	334
<i>Upcycle Your OCR: Reusing OCRs for Post-OCR Text Correction in Romanised Sanskrit</i>	
Amrith Krishna, Bodhisattwa P. Majumder, Rajesh Bhat and Pawan Goyal	345
<i>Weakly-Supervised Neural Semantic Parsing with a Generative Ranker</i>	
Jianpeng Cheng and Mirella Lapata	356
<i>Modeling Composite Labels for Neural Morphological Tagging</i>	
Alexander Tkachenko and Kairit Sirts	368
<i>Evolutionary Data Measures: Understanding the Difficulty of Text Classification Tasks</i>	
Edward Collins, Nikolai Rozanov and Bingbing Zhang	380
<i>Vectorial Semantic Spaces Do Not Encode Human Judgments of Intervention Similarity</i>	
Paola Merlo and Francesco Ackermann	392
<i>Lessons Learned in Multilingual Grounded Language Learning</i>	
Ákos Kádár, Desmond Elliott, Marc-Alexandre Côté, Grzegorz Chrupała and Afra Alishahi ..	402
<i>Unsupervised Sentence Compression using Denoising Auto-Encoders</i>	
Thibault Fevry and Jason Phang	413
<i>Resources to Examine the Quality of Word Embedding Models Trained on n-Gram Data</i>	
Ábel Elekes, Adrian Englhardt, Martin Schäler and Klemens Böhm	423
<i>Linguistically-Based Deep Unstructured Question Answering</i>	
Ahmad Aghaebrahimian	433
<i>DIMSIM: An Accurate Chinese Phonetic Similarity Algorithm Based on Learned High Dimensional Encoding</i>	
Min Li, Marina Danilevsky, Sara Noeman and Yunyao Li	444
<i>Challenge or Empower: Revisiting Argumentation Quality in a News Editorial Corpus</i>	
Roxanne El Baff, Henning Wachsmuth, Khalid Al Khatib and Benno Stein	454
<i>Bringing Order to Neural Word Embeddings with Embeddings Augmented by Random Permutations (EARP)</i>	
Trevor Cohen and Dominic Widdows	465
<i>Aggregated Semantic Matching for Short Text Entity Linking</i>	
Feng Nie, Shuyan Zhou, Jing Liu, Jinpeng Wang, Chin-Yew Lin and Rong Pan	476
<i>Adversarial Over-Sensitivity and Over-Stability Strategies for Dialogue Models</i>	
Tong Niu and Mohit Bansal	486
<i>Improving Response Selection in Multi-Turn Dialogue Systems by Incorporating Domain Knowledge</i>	
Debanjan Chaudhuri, Agustinus Kristiadi, Jens Lehmann and Asja Fischer	497
<i>The Lifted Matrix-Space Model for Semantic Composition</i>	
WooJin Chung, Sheng-Fu Wang and Samuel Bowman	508
<i>End-to-End Neural Entity Linking</i>	
Nikolaos Kolitsas, Octavian-Eugen Ganea and Thomas Hofmann	519

<i>Modelling Salient Features as Directions in Fine-Tuned Semantic Spaces</i> Thomas Ager, Ondřej Kuželka and Steven Schockaert	530
<i>Model Transfer with Explicit Knowledge of the Relation between Class Definitions</i> Hiyori Yoshikawa and Tomoya Iwakura.....	541
<i>Aiming to Know You Better Perhaps Makes Me a More Engaging Dialogue Partner</i> Yury Zemlyanskiy and Fei Sha	551
<i>Neural Maximum Subgraph Parsing for Cross-Domain Semantic Dependency Analysis</i> Yufei Chen, Sheng Huang, Fang Wang, Junjie Cao, Weiwei Sun and Xiaojun Wan	562
<i>From Random to Supervised: A Novel Dropout Mechanism Integrated with Global Information</i> Hengru Xu, Shen Li, Renfen Hu, Si Li and Sheng Gao	573
<i>Sequence to Sequence Mixture Model for Diverse Machine Translation</i> Xuanli He, Gholamreza Haffari and Mohammad Norouzi	583

Conference Program

Wednesday, October 31, 2018

08:45–09:00 *Opening remarks*

09:00–10:30 *CoNLL Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*

10:30–11:00 *Coffee break*

11:00–12:30 *CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*

12:30–14:00 *Lunch*

Invited Talk by Max Welling

14:00–15:00 *Inductive Bias in Deep Learning*

Session 1

15:00–15:15 *Embedded-State Latent Conditional Random Fields for Sequence Labeling*
Dung Thai, Sree Harsha Ramesh, Shikhar Murty, Luke Vilnis and Andrew McCallum

15:15–15:30 *Continuous Word Embedding Fusion via Spectral Decomposition*
Tianfan Fu, Cheng Zhang and Stephan Mandt

15:30–16:00 *Coffee break*

16:00–17:30 *Poster session 1*

Dual Latent Variable Model for Low-Resource Natural Language Generation in Dialogue Systems

Van-Khanh Tran and Le-Minh Nguyen

Wednesday, October 31, 2018 (continued)

A Trio Neural Model for Dynamic Entity Relatedness Ranking

Tu Nguyen, Tuan Tran and Wolfgang Nejdl

A Unified Neural Network Model for Geolocating Twitter Users

Mohammad Ebrahimi, Elaheh ShafieiBavani, Raymond Wong and Fang Chen

Corpus-Driven Thematic Hierarchy Induction

Iliia Kuznetsov and Iryna Gurevych

Adversarially Regularising Neural NLI Models to Integrate Logical Background Knowledge

Pasquale Minervini and Sebastian Riedel

From Strings to Other Things: Linking the Neighborhood and Transposition Effects in Word Reading

Stephan Tulkens, Dominiek Sandra and Walter Daelemans

Global Attention for Name Tagging

Boliang Zhang, Spencer Whitehead, Lifu Huang and Heng Ji

Pervasive Attention: 2D Convolutional Neural Networks for Sequence-to-Sequence Prediction

Maha Elbayad, Laurent Besacier and Jakob Verbeek

Comparing Attention-Based Convolutional and Recurrent Neural Networks: Success and Limitations in Machine Reading Comprehension

Matthias Blohm, Glorianna Jagfeld, Ekta Sood, Xiang Yu and Ngoc Thang Vu

Uncovering Code-Mixed Challenges: A Framework for Linguistically Driven Question Generation and Neural Based Question Answering

Deepak Gupta, Pabitra Lenka, Asif Ekbal and Pushpak Bhattacharyya

Learning to Embed Semantic Correspondence for Natural Language Understanding

Sangkeun Jung, Jinsik Lee and Jiwon Kim

Commonsense Knowledge Base Completion and Generation

Itsumi Saito, Kyosuke Nishida, Hisako Asano and Junji Tomita

Active Learning for Interactive Neural Machine Translation of Data Streams

Álvaro Peris and Francisco Casacuberta

Wednesday, October 31, 2018 (continued)

Churn Intent Detection in Multilingual Chatbot Conversations and Social Media

Christian Abbet, Meryem M'hamdi, Athanasios Giannakopoulos, Robert West, Andreea Hossmann, Michael Baeriswyl and Claudiu Musat

Learning Text Representations for 500K Classification Tasks on Named Entity Disambiguation

Ander Barrena, Aitor Soroa and Eneko Agirre

Hierarchical Attention Based Position-Aware Network for Aspect-Level Sentiment Analysis

Lishuang Li, Yang Liu and AnQiao Zhou

Bidirectional Generative Adversarial Networks for Neural Machine Translation

Zhirui Zhang, Shujie Liu, Mu Li, Ming Zhou and Enhong Chen

Latent Entities Extraction: How to Extract Entities that Do Not Appear in the Text?

Eylon Shoshan and Kira Radinsky

Generalizing Procrustes Analysis for Better Bilingual Dictionary Induction

Yova Kementchedjhieva, Sebastian Ruder, Ryan Cotterell and Anders Søgaard

Simple Unsupervised Keyphrase Extraction using Sentence Embeddings

Kamil Bennani-Smires, Claudiu Musat, Andreea Hossmann, Michael Baeriswyl and Martin Jaggi

Thursday, November 1, 2018

Session 2

09:00–09:15 *A Temporally Sensitive Submodularity Framework for Timeline Summarization*

Sebastian Martschat and Katja Markert

09:15–09:30 *Chinese Poetry Generation with a Salient-Clue Mechanism*

Xiaoyuan Yi, Ruoyu Li and Maosong Sun

09:30–09:45 *Multi-Modal Sequence Fusion via Recursive Attention for Emotion Recognition*

Rory Beard, Ritwik Das, Raymond W. M. Ng, P. G. Keerthana Gopalakrishnan, Luka Eerens, Pawel Swietojanski and Ondrej Miksik

Thursday, November 1, 2018 (continued)

09:45–10:00 *Using Sparse Semantic Embeddings Learned from Multimodal Text and Image Data to Model Human Conceptual Knowledge*

Steven Derby, Paul Miller, Brian Murphy and Barry Devereux

10:00–10:15 *Similarity Dependent Chinese Restaurant Process for Cognate Identification in Multilingual Wordlists*

Taraka Rama

10:15–10:30 *Uncovering Divergent Linguistic Information in Word Embeddings with Lessons for Intrinsic and Extrinsic Evaluation*

Mikel Artetxe, Gorka Labaka, Inigo Lopez-Gazpio and Eneko Agirre

10:30–11:00 *Coffee break*

Invited talk by Asifa Majid

11:00–12:00 *Semantic Spaces Across Diverse Languages*

Session 3

12:00–12:15 *Comparing Models of Associative Meaning: An Empirical Investigation of Reference in Simple Language Games*

Judy Hanwen Shen, Matthias Hofer, Bjarke Felbo and Roger Levy

12:15–12:30 *Sequence Classification with Human Attention*

Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei and Anders Søgaard

12:30–14:00 *Lunch*

Thursday, November 1, 2018 (continued)

Session 4

- 14:00–14:15 *Sentence-Level Fluency Evaluation: References Help, But Can Be Spared!*
Katharina Kann, Sascha Rothe and Katja Filippova
- 14:15–14:30 *Predefined Sparseness in Recurrent Sequence Models*
Thomas Demeester, Johannes Deleu, Frédéric Godin and Chris Develder
- 14:30–14:45 *Learning to Actively Learn Neural Machine Translation*
Ming Liu, Wray Buntine and Gholamreza Haffari
- 14:45–15:00 *Upcycle Your OCR: Reusing OCRs for Post-OCR Text Correction in Romanised Sanskrit*
Amrith Krishna, Bodhisattwa P. Majumder, Rajesh Bhat and Pawan Goyal
- 15:00–15:15 *Weakly-Supervised Neural Semantic Parsing with a Generative Ranker*
Jianpeng Cheng and Mirella Lapata
- 15:15–15:30 *Modeling Composite Labels for Neural Morphological Tagging*
Alexander Tkachenko and Kairit Sirts
- 15:30–16:00 *Coffee break*

16:00–17:30 Poster session 2

- Evolutionary Data Measures: Understanding the Difficulty of Text Classification Tasks*
Edward Collins, Nikolai Rozanov and Bingbing Zhang
- Vectorial Semantic Spaces Do Not Encode Human Judgments of Intervention Similarity*
Paola Merlo and Francesco Ackermann
- Lessons Learned in Multilingual Grounded Language Learning*
Ákos Kádár, Desmond Elliott, Marc-Alexandre Côté, Grzegorz Chrupała and Afra Alishahi

Thursday, November 1, 2018 (continued)

Unsupervised Sentence Compression using Denoising Auto-Encoders

Thibault Fevry and Jason Phang

Resources to Examine the Quality of Word Embedding Models Trained on n-Gram Data

Ábel Elekes, Adrian Englhardt, Martin Schäler and Klemens Böhm

Linguistically-Based Deep Unstructured Question Answering

Ahmad Aghaebrahimian

DIMSIM: An Accurate Chinese Phonetic Similarity Algorithm Based on Learned High Dimensional Encoding

Min Li, Marina Danilevsky, Sara Noeman and Yunyao Li

Challenge or Empower: Revisiting Argumentation Quality in a News Editorial Corpus

Roxanne El Baff, Henning Wachsmuth, Khalid Al Khatib and Benno Stein

Bringing Order to Neural Word Embeddings with Embeddings Augmented by Random Permutations (EARP)

Trevor Cohen and Dominic Widdows

Aggregated Semantic Matching for Short Text Entity Linking

Feng Nie, Shuyan Zhou, Jing Liu, Jinpeng Wang, Chin-Yew Lin and Rong Pan

Adversarial Over-Sensitivity and Over-Stability Strategies for Dialogue Models

Tong Niu and Mohit Bansal

Improving Response Selection in Multi-Turn Dialogue Systems by Incorporating Domain Knowledge

Debanjan Chaudhuri, Agustinus Kristiadi, Jens Lehmann and Asja Fischer

The Lifted Matrix-Space Model for Semantic Composition

WooJin Chung, Sheng-Fu Wang and Samuel Bowman

End-to-End Neural Entity Linking

Nikolaos Kolitsas, Octavian-Eugen Ganea and Thomas Hofmann

Modelling Salient Features as Directions in Fine-Tuned Semantic Spaces

Thomas Ager, Ondřej Kuželka and Steven Schockaert

Thursday, November 1, 2018 (continued)

Model Transfer with Explicit Knowledge of the Relation between Class Definitions

Hiyori Yoshikawa and Tomoya Iwakura

Aiming to Know You Better Perhaps Makes Me a More Engaging Dialogue Partner

Yury Zemlyanskiy and Fei Sha

Neural Maximum Subgraph Parsing for Cross-Domain Semantic Dependency Analysis

Yufei Chen, Sheng Huang, Fang Wang, Junjie Cao, Weiwei Sun and Xiaojun Wan

From Random to Supervised: A Novel Dropout Mechanism Integrated with Global Information

Hengru Xu, Shen Li, Renfen Hu, Si Li and Sheng Gao

Sequence to Sequence Mixture Model for Diverse Machine Translation

Xuanli He, Gholamreza Haffari and Mohammad Norouzi

Keynote Talk

Inductive Bias in Deep Learning

Max Welling

Deep learning is often considered a ‘black box’ predictor, that is, a highly flexible mapping from input variables to target variables which is hard to interpret. In almost all other scientific disciplines researchers build highly intuitive models with few variables in which decades of accumulated expertise is embedded. Not surprisingly, black box models need a lot of data to be successful as predictors while generative models need much less data. One natural question to ask is if we can inject more inductive bias in black box models, such as deep neural networks.

We will look at two different ways to achieve this. First, data often has certain symmetries, i.e. a satellite image will have no useful information in the orientation of the objects of interest. This is of course similar to the fact that in natural images there is typically no useful information in the absolute location of the objects. Convolutions implement the latter inductive bias and lead to very significant gains in terms of data efficiency. We will argue that there may be other symmetries present in data (such as orientation) which can also be hardcoded in a deep architecture and result in data efficiency gains. We will illustrate this idea in pathology slide analysis.

A second way to inject inductive bias into predictors is to consider the data generating process of the data. I will argue that for certain tasks, such as image reconstruction, the generative process can be directly embedded into the classifier by, at every layer of the network, comparing the data generated by the current reconstruction with the observations and feeding the difference back into the network. We will illustrate the resulting model, which we call the “Recurrent Inference Machine” on the task MRI image reconstruction.

Keynote Talk
Semantic Spaces Across Diverse Languages
Asifa Majid

Across diverse disciplines there is a wide-spread assumption that natural languages are equally expressible: anything that can be thought can be said. In fact, words are held to label categories that exist independently of language, such that language merely captures these pre-existing categories. In this talk, I will illustrate through cross-linguistic comparison across diverse domains that named distinctions are not nearly as self-evident as they may seem on first examination. Even for basic perceptual experiences, languages vary in which notions they lexicalise, and which concepts are coded at all. Crucially, in order to develop a universal theory of semantics, scholars must first seriously engage with the cultural variation found worldwide.