

ACL 2016

**The 54th Annual Meeting of the
Association for Computational Linguistics**

**Proceedings of the 1st Workshop on Representation Learning
for NLP**

August 11th, 2016
Berlin, Germany

©2016 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 978-1-945626-04-3

Introduction

Welcome to the 1st Workshop on Representation Learning for NLP (Repl4NLP), held on August 11, 2016 and hosted by the 54th Annual Meeting of the Association for Computational Linguistics (ACL) in Berlin, Germany. The workshop is sponsored by DeepMind, Facebook AI Research, and Microsoft Research.

Representation Learning for NLP aims to continue the spirit of previously successful workshops at ACL/NAACL/EACL, namely VSM at NAACL'15 and CVSC at ACL'13/EACL'14/ACL'15, which focussed on vector space models of meaning, compositionality, and the application of deep neural networks and spectral methods to NLP. It provides a forum for discussing recent advances on these topics, as well as future research directions in linguistically motivated vector-based models in NLP.

Organizers:

Phil Blunsom, DeepMind and Oxford University
Kyunghyun Cho, New York University
Shay Cohen, University of Edinburgh
Edward Grefenstette, DeepMind
Karl Moritz Hermann, DeepMind
Laura Rimell, University of Cambridge
Jason Weston, Facebook
Scott Yih, Microsoft

Program Committee:

Marco Baroni, University of Trento
Antoine Bordes, Facebook
Leon Bottou, Facebook
Xavier Carreras, Xerox
Stephen Clark, University of Cambridge
Hal Daumé III, University of Maryland
Kevin Duh, Johns Hopkins University
Pino Di Fabrizio, Amazon
Manaal Faruqi, Carnegie Mellon University
Dean Foster, University of Pennsylvania
Yoav Goldberg, Bar Ilan University
Jamie Ryan Kiros, Toronto
Tomáš Kočiský, DeepMind
Omer Levy, Bar Ilan University
Wang Ling, DeepMind
Graham Neubig, NAIST
Ankur Parikh, Google
John Platt, Google
Roi Reichart, Technion
Sebastian Riedel, UCL
Tim Rocktaschel, UCL
Diarmuid O Seaghdha, University of Cambridge / VocalIQ
Richard Socher, Salesforce MetaMind
Mark Steedman, University of Edinburgh
Karl Stratos, Columbia University
Peter Turney, Allen Institute for Artificial Intelligence (AI2)
Lyle Ungar, University of Pennsylvania
Oriol Vinyals, DeepMind
Guillaume Wisniewski, LIMSI-CNRS

Table of Contents

<i>Explaining Predictions of Non-Linear Classifiers in NLP</i> Leila Arras, Franziska Horn, Grégoire Montavon, Klaus-Robert Müller and Wojciech Samek	1
<i>Joint Learning of Sentence Embeddings for Relevance and Entailment</i> Petr Baudiš, Silvestr Stanko and Jan Šedivý	8
<i>A Joint Model for Word Embedding and Word Morphology</i> Kris Cao and Marek Rei	18
<i>On the Compositionality and Semantic Interpretation of English Noun Compounds</i> Corina Dima	27
<i>Functional Distributional Semantics</i> Guy Emerson and Ann Copestake	40
<i>Assisting Discussion Forum Users using Deep Recurrent Neural Networks</i> Jacob Hagstedt P Suorra and Olof Mogren	53
<i>Adjusting Word Embeddings with Semantic Intensity Orders</i> Joo-Kyung Kim, Marie-Catherine de Marneffe and Eric Fosler-Lussier	62
<i>Towards Abstraction from Extraction: Multiple Timescale Gated Recurrent Unit for Summarization</i> Minsoo Kim, Dennis Singh Moirangthem and Minh Lee	70
<i>An Empirical Evaluation of doc2vec with Practical Insights into Document Embedding Generation</i> Jey Han Lau and Timothy Baldwin	78
<i>Quantifying the Vanishing Gradient and Long Distance Dependency Problem in Recursive Neural Networks and Recursive LSTMs</i> Phong Le and Willem Zuidema	87
<i>LSTM-Based Mixture-of-Experts for Knowledge-Aware Dialogues</i> Phong Le, Marc Dymetman and Jean-Michel Renders	94
<i>Mapping Unseen Words to Task-Trained Embedding Spaces</i> Pranava Swaroop Madhyastha, Mohit Bansal, Kevin Gimpel and Karen Livescu	100
<i>Multilingual Modal Sense Classification using a Convolutional Neural Network</i> Ana Marasović and Anette Frank	111
<i>Towards cross-lingual distributed representations without parallel text trained with adversarial autoencoders</i> Antonio Valerio Miceli Barone	121
<i>Decomposing Bilexical Dependencies into Semantic and Syntactic Vectors</i> Jeff Mitchell	127
<i>Learning Semantic Relatedness in Community Question Answering Using Neural Models</i> Henry Nassif, Mitra Mohtarami and James Glass	137
<i>Learning Text Similarity with Siamese Recurrent Networks</i> Paul Neculoiu, Maarten Versteegh and Mihai Rotaru	148

<i>A Two-stage Approach for Extending Event Detection to New Types via Neural Networks</i>	
Thien Huu Nguyen, Lisheng Fu, Kyunghyun Cho and Ralph Grishman	158
<i>Parameterized context windows in Random Indexing</i>	
Tobias Norlund, David Nilsson and Magnus Sahlgren	166
<i>Making Sense of Word Embeddings</i>	
Maria Pelevina, Nikolay Arefiev, Chris Biemann and Alexander Panchenko	174
<i>Pair Distance Distribution: A Model of Semantic Representation</i>	
Yonatan Ramni, Oded Maimon and Evgeni Khmelnitsky	184
<i>Measuring Semantic Similarity of Words Using Concept Networks</i>	
Gábor Recski, Eszter Iklódi, Katalin Pajkossy and Andras Kornai	193
<i>Using Embedding Masks for Word Categorization</i>	
Stefan Ruseti, Traian Rebedea and Stefan Trausan-Matu	201
<i>Sparsifying Word Representations for Deep Unordered Sentence Modeling</i>	
Prasanna Sattigeri and Jayaraman J. Thiagarajan	206
<i>Why "Blow Out"? A Structural Analysis of the Movie Dialog Dataset</i>	
Richard Searle and Megan Bingham-Walker	215
<i>Learning Word Importance with the Neural Bag-of-Words Model</i>	
Imran Sheikh, Irina Illina, Dominique Fohr and Georges Linarès	222
<i>A Vector Model for Type-Theoretical Semantics</i>	
Konstantin Sokolov	230
<i>Towards Generalizable Sentence Embeddings</i>	
Eleni Triantafillou, Jamie Ryan Kiros, Raquel Urtasun and Richard Zemel	239
<i>Domain Adaptation for Neural Networks by Parameter Augmentation</i>	
Yusuke Watanabe, Kazuma Hashimoto and Yoshimasa Tsuruoka	249
<i>Neural Associative Memory for Dual-Sequence Modeling</i>	
Dirk Weissenborn	258

Conference Program

Thursday, August 11, 2016

9:30–9:40 Welcome and Opening Remarks

9:40–10:30 Keynote: Katrin Erk (University of Texas at Austin)

10:30–11:00 Coffee Break

11:00–11:50 Keynote: Animashree Anandkumar (University of California, Irvine)

11:50–12:10 Best Papers

12:10–13:30 Lunch Break

13:30–14:20 Keynote: Hal Daumé III (University of Maryland)

14:20–15:10 Keynote: Raia Hadsell (DeepMind)

15:10–15:30 Poster Session

Decoding Neural Activity Patterns Associated with Sentences by Combining Experiential Attribute and Text-Based Semantic Models

Andrew Anderson, Jeffrey Binder, Leonardo Fernandino, Colin Humphries, Lisa Conant, Katrin Erk and Rajeev Raizada

Explaining Predictions of Non-Linear Classifiers in NLP

Leila Arras, Franziska Horn, Grégoire Montavon, Klaus-Robert Müller and Wojciech Samek

Joint Learning of Sentence Embeddings for Relevance and Entailment

Petr Baudiš, Silvestr Stanko and Jan Šedivý

Combining String Kernels and Gaussian Processes for Richer Text Representations.

Daniel Beck

Thursday, August 11, 2016 (continued)

A Joint Model for Word Embedding and Word Morphology

Kris Cao and Marek Rei

Learning Word Representations from Multiple Information Sources

Yunchuan Chen, Lili Mou, Yan Xu, Ge Li and Zhi Jin

On the Compositionality and Semantic Interpretation of English Noun Compounds

Corina Dima

Functional Distributional Semantics

Guy Emerson and Ann Copestake

Improving Preposition Sense Disambiguation with Representations Learned from Multilingual Data

Hila Gonen and Yoav Goldberg

Assisting Discussion Forum Users using Deep Recurrent Neural Networks

Jacob Hagstedt P Suorra and Olof Mogren

Adjusting Word Embeddings with Semantic Intensity Orders

Joo-Kyung Kim, Marie-Catherine de Marneffe and Eric Fosler-Lussier

Towards Abstraction from Extraction: Multiple Timescale Gated Recurrent Unit for Summarization

Minsoo Kim, Dennis Singh Moirangthem and Minhoo Lee

An Empirical Evaluation of doc2vec with Practical Insights into Document Embedding Generation

Jey Han Lau and Timothy Baldwin

Quantifying the Vanishing Gradient and Long Distance Dependency Problem in Recursive Neural Networks and Recursive LSTMs

Phong Le and Willem Zuidema

LSTM-Based Mixture-of-Experts for Knowledge-Aware Dialogues

Phong Le, Marc Dymetman and Jean-Michel Renders

Learning Phone Embeddings for Word Segmentation of Child-Directed Speech

Jianqiang Ma, Çağrı Çöltekin and Erhard Hinrichs

Mapping Unseen Words to Task-Trained Embedding Spaces

Pranava Swaroop Madhyastha, Mohit Bansal, Kevin Gimpel and Karen Livescu

Thursday, August 11, 2016 (continued)

Multilingual Modal Sense Classification using a Convolutional Neural Network

Ana Marasović and Anette Frank

Towards cross-lingual distributed representations without parallel text trained with adversarial autoencoders

Antonio Valerio Miceli Barone

Decomposing Bilexical Dependencies into Semantic and Syntactic Vectors

Jeff Mitchell

Distilling Word Embeddings: An Encoding Approach

Lili Mou, Ran Jia, Yan Xu, Ge Li, Lu Zhang and Zhi Jin

Learning Semantic Relatedness in Community Question Answering Using Neural Models

Henry Nassif, Mitra Mohtarami and James Glass

Learning Text Similarity with Siamese Recurrent Networks

Paul Neculoiu, Maarten Versteegh and Mihai Rotaru

A Two-stage Approach for Extending Event Detection to New Types via Neural Networks

Thien Huu Nguyen, Lisheng Fu, Kyunghyun Cho and Ralph Grishman

Parameterized context windows in Random Indexing

Tobias Norlund, David Nilsson and Magnus Sahlgren

Making Sense of Word Embeddings

Maria Pelevina, Nikolay Arefiev, Chris Biemann and Alexander Panchenko

Pair Distance Distribution: A Model of Semantic Representation

Yonatan Ramni, Oded Maimon and Evgeni Khmelnitsky

Measuring Semantic Similarity of Words Using Concept Networks

Gábor Recski, Eszter Iklódi, Katalin Pajkossy and Andras Kornai

Using Embedding Masks for Word Categorization

Stefan Ruseti, Traian Rebedea and Stefan Trausan-Matu

Sparsifying Word Representations for Deep Unordered Sentence Modeling

Prasanna Sattigeri and Jayaraman J. Thiagarajan

Thursday, August 11, 2016 (continued)

Why "Blow Out"? A Structural Analysis of the Movie Dialog Dataset

Richard Searle and Megan Bingham-Walker

Learning Word Importance with the Neural Bag-of-Words Model

Imran Sheikh, Irina Illina, Dominique Fohr and Georges Linarès

A Vector Model for Type-Theoretical Semantics

Konstantin Sokolov

Towards Generalizable Sentence Embeddings

Eleni Triantafillou, Jamie Ryan Kiros, Raquel Urtasun and Richard Zemel

Domain Adaptation for Neural Networks by Parameter Augmentation

Yusuke Watanabe, Kazuma Hashimoto and Yoshimasa Tsuruoka

Neural Associative Memory for Dual-Sequence Modeling

Dirk Weissenborn

MuFuRU: The Multi-Function Recurrent Unit

Dirk Weissenborn and Tim Rocktäschel

15.30–16.00 Poster Session Continues and Coffee break

16.00–17.20 Panel Discussion

Thursday, August 11, 2016 (continued)

17.20–17.30 Closing Remarks

