

ACL 2019

**The 4th Workshop on Representation Learning for NLP
(RepL4NLP-2019)**

Proceedings of the Workshop

August 2, 2019
Florence, Italy

Sponsors:

facebook

amazon

LA33
NAVER LABS
EUROPE

©2019 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 978-1-950737-35-2

Introduction

The 4th Workshop on Representation Learning for NLP (RepL4NLP) will be hosted by ACL 2019 and held on 2 August 2019. The workshop is being organised by Isabelle Augenstein, Spandana Gella, Sebastian Ruder, Katharina Kann, Burcu Can, Alexis Conneau, Johannes Welbl, Xian Ren and Marek Rei; and advised by Kyunghyun Cho, Edward Grefenstette, Karl Moritz Hermann, Chris Dyer and Laura Rimell. The workshop is organised by the ACL Special Interest Group on Representation Learning (SIGREP) and receives generous sponsorship from Facebook AI Research, Amazon, and Naver.

The 4th Workshop on Representation Learning for NLP aims to continue the success of the 1st Workshop on Representation Learning for NLP (about 50 submissions and over 250 attendees; second most attended collocated event at ACL'16 after WMT), 2nd Workshop on Representation Learning for NLP and 3rd Workshop on Representation Learning for NLP. The workshop was introduced as a synthesis of several years of independent *CL workshops focusing 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:

Isabelle Augenstein, University of Copenhagen
Burcu Can, Hacettepe University
Alexis Conneau, Facebook AI Research Paris, Université Le Mans
Spandana Gella, Amazon AI
Katharina Kann, New York University
Marek Rei, University of Cambridge
Xiang Ren, University of Southern California
Sebastian Ruder, DeepMind
Johannes Welbl, University College London

Senior Advisors:

Kyunghyun Cho, NYU and Facebook AI Research
Chris Dyer, DeepMind
Edward Grefenstette, Facebook AI Research
Karl Moritz Hermann, DeepMind
Laura Rimell, DeepMind

Program Committee:

Eneko Agirre, University of the Basque Country (UPV/EHU)
Yoav Artzi, Cornell University
Jonathan Berant, Tel Aviv University and AI2
Johannes Bjerva, Department of Computer Science, University of Copenhagen
Samuel R. Bowman, New York University
Jan Buys, University of Washington
Andrew Caines, University of Cambridge
Claire Cardie, Cornell University
Xilun Chen, Cornell University
Danqi Chen, Stanford University
Heeyoul Choi, Handong Global University
Eunsol Choi, University of Washington
Manuel Ciosici, Aarhus University
Stephen Clark, DeepMind
Marco Damonte, The University of Edinburgh
Miryam de Lhoneux, Uppsala University
Desmond Elliott, University of Copenhagen
Orhan Firat, Google AI
Lucie Flekova, Amazon Research
Daniela Gerz, University of Cambridge
Kevin Gimpel, Toyota Technological Institute at Chicago
Gholamreza Haffari, Monash University
Mareike Hartmann, University of Copenhagen
Mohit Iyyer, University of Massachusetts Amherst

Arzoo Katiyar, Cornell University
Yova Kementchedjhieva, University of Copenhagen
Douwe Kiela, Facebook
Bill Yuchen Lin, USC/ISI
Suresh Manandhar, University of York
Ana Marasovic, Allen Institute for Artificial Intelligence (AI2)
Sebastian J. Mielke, Johns Hopkins University
Todor Mihaylov, Heidelberg University
Pasquale Minervini, UCL
Nikita Nangia, New York University
Jason Naradowsky, Johns Hopkins
Shashi Narayan, Google
Thien Huu Nguyen, University of Oregon
Inkit Padhi, IBM Research
Chris Quirk, Microsoft Research
Roi Reichart, Technion - Israel Institute of Technology
Alan Ritter, The Ohio State University
Hinrich Schütze, Center for Information and Language Processing, University of Munich
Minjoon Seo, University of Washington
Tianze Shi, Cornell University
Vered Shwartz, Bar-Ilan University
Daniil Sorokin, UKP Lab, Technische Universität Darmstadt
Mark Steedman, University of Edinburgh
Karl Stratos, Toyota Technological Institute at Chicago
Jörg Tiedemann, University of Helsinki
Ivan Titov, University of Edinburgh / University of Amsterdam
Eva Maria Vecchi, University of Cambridge
Ivan Vulić, University of Cambridge
Dirk Weissenborn, Google AI, Berlin
Yadollah Yaghoobzadeh, Microsoft Research Montreal
Yi Yang, ASAPP
Helen Yannakoudakis, University of Cambridge
Wenpeng Yin, University of Pennsylvania
Luke Zettlemoyer, University of Washington
Wenxuan Zhou, University of Southern California
Diarmuid Ó Séaghdha, Apple
Robert Östling, Department of Linguistics, Stockholm University

Keynote Speakers:

Mohit Bansal, UNC Chapel Hill
Marco Baroni, Facebook AI Research
Raquel Fernandez, University of Amsterdam
Yulia Tsvetkov, Carnegie Mellon University

Table of Contents

<i>Deep Generalized Canonical Correlation Analysis</i> Adrian Benton, Huda Khayrallah, Biman Gujral, Dee Ann Reisinger, Sheng Zhang and Raman Arora.....	1
<i>To Tune or Not to Tune? Adapting Pretrained Representations to Diverse Tasks</i> Matthew E. Peters, Sebastian Ruder and Noah A. Smith	7
<i>Generative Adversarial Networks for Text Using Word2vec Intermediaries</i> Akshay Budhkar, Krishnapriya Vishnubhotla, Safwan Hossain and Frank Rudzicz.....	15
<i>An Evaluation of Language-Agnostic Inner-Attention-Based Representations in Machine Translation</i> Alessandro Raganato, Raúl Vázquez, Mathias Creutz and Jörg Tiedemann.....	27
<i>Multilingual NMT with a Language-Independent Attention Bridge</i> Raúl Vázquez, Alessandro Raganato, Jörg Tiedemann and Mathias Creutz.....	33
<i>Efficient Language Modeling with Automatic Relevance Determination in Recurrent Neural Networks</i> Maxim Kodryan, Artem Grachev, Dmitry Ignatov and Dmitry Vetrov	40
<i>MoRTy: Unsupervised Learning of Task-specialized Word Embeddings by Autoencoding</i> Nils Rethmeier and Barbara Plank.....	49
<i>Pitfalls in the Evaluation of Sentence Embeddings</i> Steffen Eger, Andreas Rücklé and Iryna Gurevych	55
<i>Learning Bilingual Sentence Embeddings via Autoencoding and Computing Similarities with a Multi-layer Perceptron</i> Yunsu Kim, Hendrik Rosendahl, Nick Rossenbach, Jan Rosendahl, Shahram Khadivi and Hermann Ney	61
<i>Specializing Distributional Vectors of All Words for Lexical Entailment</i> Aishwarya Kamath, Jonas Pfeiffer, Edoardo Maria Ponti, Goran Glavaš and Ivan Vulić	72
<i>Composing Noun Phrase Vector Representations</i> Aikaterini-Lida Kalouli, Valeria dePaiva and Richard Crouch	84
<i>Towards Robust Named Entity Recognition for Historic German</i> Stefan Schweter and Johannes Baiter	96
<i>On Evaluating Embedding Models for Knowledge Base Completion</i> Yanjie Wang, Daniel Ruffinelli, Rainer Gemulla, Samuel Broscheit and Christian Meilicke....	104
<i>Constructive Type-Logical Supertagging With Self-Attention Networks</i> Konstantinos Kogkalidis, Michael Moortgat and Tejaswini Deoskar.....	113
<i>Auto-Encoding Variational Neural Machine Translation</i> Bryan Eikema and Wilker Aziz	124
<i>Learning Bilingual Word Embeddings Using Lexical Definitions</i> Weijia Shi, Muhao Chen, Yingtao Tian and Kai-Wei Chang	142
<i>An Empirical Study on Pre-trained Embeddings and Language Models for Bot Detection</i> Andres Garcia-Silva, Cristian Berrio and José Manuel Gómez-Pérez.....	148

<i>Probing Multilingual Sentence Representations With X-Probe</i> Vinit Ravishankar, Lilja Øvrelid and Erik Velldal	156
<i>Fine-Grained Entity Typing in Hyperbolic Space</i> Federico López, Benjamin Heinzerling and Michael Strube	169
<i>Learning Multilingual Meta-Embeddings for Code-Switching Named Entity Recognition</i> Genta Indra Winata, Zhaojiang Lin and Pascale Fung	181
<i>Investigating Sub-Word Embedding Strategies for the Morphologically Rich and Free Phrase-Order Hungarian</i> Bálint Döbrössi, Márton Makrai, Balázs Tarján and György Szaszák	187
<i>A Self-Training Approach for Short Text Clustering</i> Amir Hadifar, Lucas Sterckx, Thomas Demeester and Chris Develder	194
<i>Improving Word Embeddings Using Kernel PCA</i> Vishwani Gupta, Sven Giesselbach, Stefan Rüping and Christian Bauckhage	200
<i>Assessing Incrementality in Sequence-to-Sequence Models</i> Dennis Ulmer, Dieuwke Hupkes and Elia Bruni	209
<i>On Committee Representations of Adversarial Learning Models for Question-Answer Ranking</i> Sparsh Gupta and Vitor Carvalho	218
<i>Meta-Learning Improves Lifelong Relation Extraction</i> Abiola Obamuyide and Andreas Vlachos	224
<i>Best Practices for Learning Domain-Specific Cross-Lingual Embeddings</i> Lena Shakurova, Beata Nyari, Chao Li and Mihai Rotaru	230
<i>Effective Dimensionality Reduction for Word Embeddings</i> Vikas Raunak, Vivek Gupta and Florian Metze	235
<i>Learning Word Embeddings without Context Vectors</i> Alexey Zobnin and Evgenia Elistratova	244
<i>Learning Cross-Lingual Sentence Representations via a Multi-task Dual-Encoder Model</i> Muthu Chidambaram, Yinfei Yang, Daniel Cer, Steve Yuan, Yunhsuan Sung, Brian Strope and Ray Kurzweil	250
<i>Modality-based Factorization for Multimodal Fusion</i> Elham J. Barezi and Pascale Fung	260
<i>Leveraging Pre-Trained Embeddings for Welsh Taggers</i> Ignatius Ezeani, Scott Piao, Steven Neale, Paul Rayson and Dawn Knight	270

Workshop Program

Friday, August 2, 2019

09:30–09:45 **Welcome and Opening Remarks**

09:45–14:45 **Keynote Session**

09:45–10:30 *Invited Talk 1*
 Marco Baroni

10:30–11:00 **Coffee Break**

11:00–11:45 *Invited Talk 2*
 Mohit Bansal

11:45–12:30 *Invited Talk 3*
 Raquel Fernández

12:30–14:00 **Lunch**

14:00–14:45 *Invited Talk 4*
 Yulia Tsvetkov

14:45–15:00 **Outstanding Papers Spotlight Presentations**

Friday, August 2, 2019 (continued)

15:00–16:30 Poster Session (including Coffee Break from 15:30-16:00) + Drinks Reception

Deep Generalized Canonical Correlation Analysis

Adrian Benton, Huda Khayrallah, Biman Gujral, Dee Ann Reisinger, Sheng Zhang and Raman Arora

CBOW Is Not All You Need: Combining CBOW with the Compositional Matrix Space Model

Florian Mai, Lukas Galke and Ansgar Scherp

To Tune or Not to Tune? Adapting Pretrained Representations to Diverse Tasks

Matthew E. Peters, Sebastian Ruder and Noah A. Smith

Generative Adversarial Networks for Text Using Word2vec Intermediaries

Akshay Budhkar, Krishnapriya Vishnubhotla, Safwan Hossain and Frank Rudzicz

An Evaluation of Language-Agnostic Inner-Attention-Based Representations in Machine Translation

Alessandro Raganato, Raúl Vázquez, Mathias Creutz and Jörg Tiedemann

Multilingual NMT with a Language-Independent Attention Bridge

Raúl Vázquez, Alessandro Raganato, Jörg Tiedemann and Mathias Creutz

Efficient Language Modeling with Automatic Relevance Determination in Recurrent Neural Networks

Maxim Kodryan, Artem Grachev, Dmitry Ignatov and Dmitry Vetrov

MoRTy: Unsupervised Learning of Task-specialized Word Embeddings by Autoencoding

Nils Rethmeier and Barbara Plank

Pitfalls in the Evaluation of Sentence Embeddings

Steffen Eger, Andreas Rücklé and Iryna Gurevych

Learning Bilingual Sentence Embeddings via Autoencoding and Computing Similarities with a Multilayer Perceptron

Yunsu Kim, Hendrik Rosendahl, Nick Rossenbach, Jan Rosendahl, Shahram Khadivi and Hermann Ney

Specializing Distributional Vectors of All Words for Lexical Entailment

Aishwarya Kamath, Jonas Pfeiffer, Edoardo Maria Ponti, Goran Glavaš and Ivan Vulić

Friday, August 2, 2019 (continued)

Composing Noun Phrase Vector Representations

Aikaterini-Lida Kalouli, Valeria dePaiva and Richard Crouch

Towards Robust Named Entity Recognition for Historic German

Stefan Schweter and Johannes Baiter

On Evaluating Embedding Models for Knowledge Base Completion

Yanjie Wang, Daniel Ruffinelli, Rainer Gemulla, Samuel Broscheit and Christian Meilicke

Constructive Type-Logical Supertagging With Self-Attention Networks

Konstantinos Kogkalidis, Michael Moortgat and Tejaswini Deoskar

Auto-Encoding Variational Neural Machine Translation

Bryan Eikema and Wilker Aziz

Responsible Team Players Wanted: An Analysis of Soft Skill Requirements in Job Advertisements

Federica Calanca, Luiza Sayfullina, Lara Minkus, Claudia Wagner and Eric Malmi

Learning Bilingual Word Embeddings Using Lexical Definitions

Weijia Shi, Muhao Chen, Yingtao Tian and Kai-Wei Chang

An Empirical Study on Pre-trained Embeddings and Language Models for Bot Detection

Andres Garcia-Silva, Cristian Berrio and José Manuel Gómez-Pérez

Probing Multilingual Sentence Representations With X-Probe

Vinit Ravishankar, Lilja Øvrelid and Erik Velldal

Fine-Grained Entity Typing in Hyperbolic Space

Federico López, Benjamin Heinzerling and Michael Strube

Semantic Cross-lingual Sentence Embeddings

Wataru Hirota, Yoshihiko Suhara, Behzad Golshan and Wang-Chiew Tan

Learning Multilingual Meta-Embeddings for Code-Switching Named Entity Recognition

Genta Indra Winata, Zhaojiang Lin and Pascale Fung

Friday, August 2, 2019 (continued)

Investigating Sub-Word Embedding Strategies for the Morphologically Rich and Free Phrase-Order Hungarian

Bálint Döbrössi, Márton Makrai, Balázs Tarján and György Szaszák

A Self-Training Approach for Short Text Clustering

Amir Hadifar, Lucas Sterckx, Thomas Demeester and Chris Develder

Extended Abstract: An Auto-NLP Representation Learning Framework

Akshay Budhkar

Improving Word Embeddings Using Kernel PCA

Vishwani Gupta, Sven Giesselbach, Stefan Rüping and Christian Bauckhage

Assessing Incrementality in Sequence-to-Sequence Models

Dennis Ulmer, Dieuwke Hupkes and Elia Bruni

On Committee Representations of Adversarial Learning Models for Question-Answer Ranking

Sparsh Gupta and Vitor Carvalho

Meta-Learning Improves Lifelong Relation Extraction

Abiola Obamuyide and Andreas Vlachos

Best Practices for Learning Domain-Specific Cross-Lingual Embeddings

Lena Shakurova, Beata Nyari, Chao Li and Mihai Rotaru

SuperTML: Two-Dimensional Word Embedding for the Precognition on Structured Tabular Data

Baohua Sun, Lin Yang, Wenhan Zhang, Michael Lin, Patrick Dong, Charles Young and Jason Dong

Effective Dimensionality Reduction for Word Embeddings

Vikas Raunak, Vivek Gupta and Florian Metze

Squared English Word: A Method of Generating Glyph to Use Super Characters for Sentiment Analysis

Baohua Sun, Lin Yang, Catherine Chi, Wenhan Zhang and Michael Lin

Learning Word Embeddings without Context Vectors

Alexey Zobnin and Evgenia Elistratova

Friday, August 2, 2019 (continued)

A Study of State Aliasing in Structured Prediction with RNNs

Layla El Asri and Adam Trischler

Learning Cross-Lingual Sentence Representations via a Multi-task Dual-Encoder Model

Muthu Chidambaram, Yinfei Yang, Daniel Cer, Steve Yuan, Yunhsuan Sung, Brian Strope and Ray Kurzweil

Modality-based Factorization for Multimodal Fusion

Elham J. Barezi and Pascale Fung

Leveraging Pre-Trained Embeddings for Welsh Taggers

Ignatius Ezeani, Scott Piao, Steven Neale, Paul Rayson and Dawn Knight

16:30–17:30 Panel Discussion

17:30–17:40 Closing Remarks + Best Paper Awards Announcement

