ACL 2023

The 8th Workshop on Representation Learning for NLP (RepL4NLP 2023)

Proceedings of the Workshop

July 13, 2023

The ACL organizers gratefully acknowledge the support from the following sponsors.

Gold

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 978-1-959429-77-7

Introduction

The 8th Workshop on Representation Learning for NLP (RepL4NLP 2023) will be hosted by ACL 2023 and held on 13 July 2023. The workshop is being organised by Burcu Can, Maximilian Mozes, Samuel Cahyawijaya, Naomi Saphra, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, and Chen Zhao; and advised by Isabelle Augenstein, Anna Rogers, Kyunghyun Cho, and Edward Grefenstette. The workshop is organised by the ACL Special Interest Group on Representation Learning (SIGREP).

The 8th Workshop on Representation Learning for NLP aims to continue the success of the Repl4NLP workshop series, with the 1st Workshop on Representation Learning for NLP having received about 50 submissions and over 250 attendees - the second most attended collocated event at ACL'16 after WMT. 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. The workshop will take place in a hybrid setting, and, as in previous years, feature interdisciplinary keynotes, paper presentations, posters, as well as a panel discussion.

Organizing Committee

Workshop Organizers

Burcu Can, University of Stirling
Maximilian Mozes, University College London
Samuel Cahyawijaya, Hong Kong University of Science and Technology
Naomi Saphra, New York University
Nora Kassner, Meta
Shauli Ravfogel, Bar-Ilan University
Abhilasha Ravichander, Allen Institute for Artificial Intelligence
Chen Zhao, New York University

Senior Advisors

Isabelle Augenstein, University of Copenhagen
Anna Rogers, University of Copenhagen
Kyunghyun Cho, New York University
Edward Grefenstette, DeepMind
Lena Voita, Meta AI

Program Committee

Program Committee

Aleksandr Drozd, RIKEN
Aleksandra Piktus, University of Roma La Sapienza
Alexandra Chronopoulou, Ludwig-Maximilians-Universität München
Alexandre Allauzen, École Supérieure de Physique et de Chimie Industrielles de la Ville de Paris
Ankur Padia, Philips Research North America
Anna Tigunova, Saarland Informatics Campus, Max-Planck Institute
Antonio Uva, Amazon
Arjun Reddy Akula, Google
Bryan Wilie, Hong Kong University of Science and Technology
Caglar Gulcehre, DeepMind
Chris Quirk, Microsoft
Christopher Davis, University of Cambridge
Denis Filimonov, Amazon
Dong Zhou, Guangdong University of Foreign Studies
Edwin Simpson, University of Bristol
Eliezer de Souza da Silva, Norwegian University of Science and Technology
Elisabetta Fersini, University of Milan-Bicocca
Etsuko Ishii, The Hong Kong University of Science and Technology
Eva Maria Vecchi, University of Stuttgart
Faeze Brahman, Allen Institute for AI
Frank Rudzicz, Dalhousie University
Hai Wang, Samsung
Haiqin Yang, International Digital Economy Academy
Hao Tang, University of Edinburgh
Holy Lovenia, Hong Kong University of Science and Technology
Hong Wang, UC Santa Barbara
Hong Yu, Apple
Hongliang Dai, The Hong Kong University of Science and Technology
Izzeddin Gur, Google
Jiaji Huang, AWS
Jingjing Xu, Shanghai AI Lab
Karl Stratos, Rutgers University
Kartik Goyal, Toyota Technological Institute at Chicago
Kimberly Mai, University College London
Lan Du, Monash University
Lili Mou, University of Alberta
Lin Chen, University of Illinois at Chicago
Linqing Liu, University College London
Matthieu Labeau, Télécom ParisTech
Menno van Zaanen, North-West University
Minhao Cheng, Hong Kong University of Science and Technology
Muhammad Mahbubur Rahman, George Washington University
Nadi Tomeh, Université Sorbonne Paris Nord
Necva Bölücü, Hacettepe University
Negar Foroutan, EPFL Lausanne
Nicholas Andrews, Johns Hopkins University

Paula Czarnowska, Amazon AWS
Philipp Dufter, Apple
Phoebe Mulcaire, Duolingo
Rainer Gemulla, Universität Mannheim
Robin Jia, University of Southern California
Rodrigo Wilkens, UCL
Santosh Kesiraju, Brno University of Technology
Selçuk Köprü, eBay
Sergey Feldman, Allen Institute for Artificial Intelligence
Shankar Kumar, Google
Shaohua Li, A*STAR
Shehzaad Zuzar Dhuliawala, Swiss Federal Institute of Technology
Shiva Taslimipoor, University of Cambridge
Shuai Tang, Amazon Web Services
Surangika Ranathunga, University of Moratuwa
Tao Li, Google
Tharindu Ranasinghe, Aston University
Thien Huu Nguyen, University of Oregon
Tiezheng YU, The Hong Kong University of Science and Technology
Tom Lippincott, Whiting School of Engineering
Truc-Vien T. Nguyen, Silo AI
Tsuyoshi Okita, Kyushu Institute of Technology
Tsvetomila Mihaylova, Aalto University
Uri Alon, Carnegie Mellon University
Valentin Hofmann, University of Oxford
Vipul Raheja, Grammarly
Vladimir Eidelman, FiscalNote, Inc.
Wenliang Dai, The Hong Kong University of Science and Technology
Wenxuan Zhou, University of Southern California
Xia Cui, The Manchester Metropolitan University
Yitong Li, Huawei Technologies Co., Ltd.
Yogarshi Vyas, Amazon
Yue Chen, Microsoft
Yuval Pinter, Ben Gurion University of the Negev

Invited Speakers

Swabha Swayamdipta, University of Southern California
Samira Abnar, Apple
Hannaneh Hajishirzi, University of Washington & Allen AI
Omer Levy, Tel Aviv University & Meta AI

Table of Contents

Adversarial Clean Label Backdoor Attacks and Defenses on Text Classification Systems
Ashim Gupta and Amrith Krishna 1

Do not Mask Randomly: Effective Domain-adaptive Pre-training by Masking In-domain Keywords
Shahriar Golchin, Mihai Surdeanu, Nazgol Tavabi and Ata Kiapour 13

Grammatical information in BERT sentence embeddings as two-dimensional arrays
Vivi Nastase and Paola Merlo 22

A Multilingual Evaluation of NER Robustness to Adversarial Inputs
Akshay Srinivasan and Sowmya Vajjala 40

Retrieval-Augmented Domain Adaptation of Language Models
Benfeng Xu, Chunxu Zhao, Wenbin Jiang, PengFei Zhu, Songtai Dai, Chao Pang, Zhuo Sun, Shuohuan Wang and Yu Sun 54

Fine-grained Text Style Transfer with Diffusion-Based Language Models
Yiwei Lyu, Tiange Luo, Jiacheng Shi, Todd C Hollon and Honglak Lee 65

Enhancing text comprehension for Question Answering with Contrastive Learning
Seungyeon Lee and Minho Lee 75

Towards Flow Graph Prediction of Open-Domain Procedural Texts
Keisuke Shirai, Hirotaka Kameko and Shinsuke Mori 87

One does not fit all! On the Complementarity of Vision Encoders for Vision and Language Tasks
Gregor Geigle, Chen Cecilia Liu, Jonas Pfeiffer and Iryna Gurevych 97

SPC: Soft Prompt Construction for Cross Domain Generalization
Wenbo Zhao, Arpit Gupta, Tagyoung Chung and Jing Huang 118

Friendly Neighbors: Contextualized Sequence-to-Sequence Link Prediction
Adrian Kochsiek, Apoorv Umang Saxena, Inderjeet Jayakumar Nair and Rainer Gemulla 131

Extracting Multi-valued Relations from Language Models
Sneha Singhania, Simon Razniewski and Gerhard Weikum 139

Hierarchical Multi-Instance Multi-Label Learning for Detecting Propaganda Techniques
Anni Chen and Bhuwan Dhingra 155

Contrastive Loss is All You Need to Recover Analogies as Parallel Lines
Narutatsu Ri, Fei-Tzin Lee and Nakul Verma 164

Syntax-Aware Graph-to-Graph Transformer for Semantic Role Labelling
Alireza Mohammadshahi and James Henderson 174

Improving Zero-shot Relation Classification via Automatically-acquired Entailment Templates
Mahdi Rahimi and Mihai Surdeanu 187

MUX-PLMs: Pre-training Language Models with Data Multiplexing
Vishvak Murahari, Ameet Deshpande, Carlos E Jimenez, Izhak Shafran, Mingqiu Wang, Yuan Cao and Karthik R Narasimhan 196

Mixed Orthographic/Phonemic Language Modeling: Beyond Orthographically Restricted Transformers (BORT)
Robert C Gale, Alexandra C Salem, Gerasimos Fergadiotis and Steven Bedrick 212

Effectiveness of Data Augmentation for Parameter Efficient Tuning with Limited Data
Stephen Obadinma, Hongyu Guo and Xiaodan Zhu 226

Relational Sentence Embedding for Flexible Semantic Matching
Bin Wang and Haizhou Li 238

Tucker Decomposition with Frequency Attention for Temporal Knowledge Graph Completion
Likang Xiao, Richong Zhang, Zijie Chen and Junfan Chen 253

CLIP-based image captioning via unsupervised cycle-consistency in the latent space
Romain Bielawski and Rufin VanRullen 266

Token-level Fitting Issues of Seq2seq Models
Guangsheng Bao, Zhiyang Teng and Yue Zhang 276

Revealing the Blind Spot of Sentence Encoder Evaluation by HEROS
Cheng-Han Chiang, Hung-yi Lee, Yung-Sung Chuang and James Glass 289

One-Shot Exemplification Modeling via Latent Sense Representations
John Harvill, Mark Hasegawa-Johnson, Hee Suk Yoon, Chang D. Yoo and Eunseop Yoon . . . 303

Sen2Pro: A Probabilistic Perspective to Sentence Embedding from Pre-trained Language Model
Lingfeng Shen, Haiyun Jiang, Lemao Liu and Shuming Shi 315

Visual Coherence Loss for Coherent and Visually Grounded Story Generation
Xudong Hong, Vera Demberg, Asad Sayeed, Qiankun Zheng and Bernt Schiele 334