SustaiNLP 2023

The Fourth Workshop on Simple and Efficient Natural Language Processing

Proceedings of the Workshop

July 13, 2023

©2023 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 978-1-959429-79-1

Introduction

It is our great pleasure to welcome you to the fourth edition of SustaiNLP: Workshop on Simple and Efficient Natural Language Processing.

The Natural Language Processing community has, in recent years, demonstrated a notable focus on improving higher scores on standard benchmarks and taking the lead on community-wide leaderboards (e.g., GLUE, SentEval). While this aspiration has led to improvements in benchmark performance of (predominantly neural) models, it has also came at a cost, i.e., increased model complexity and the evergrowing amount of computational resources required for training and using the current state-of-the-art models. Moreover, the recent research efforts have, for the most part, failed to identify sources of empirical gains in models, often failing to empirically justify the model complexity beyond benchmark performance.

Because of these easily observable trends, we organized the SustaiNLP workshop with the goal of promoting more sustainable NLP research and practices, with two main objectives: (1) encouraging development of more efficient NLP models; and (2) providing simpler architectures and empirical justification of model complexity. For both aspects, we encouraged submissions from all topical areas of NLP.

This year, we received 46 submissions, proposing a multitude of viable resource-efficient NLP methods and spanning a wide range of NLP applications. We have selected 26 submissions for presentation at the workshop, yielding an acceptance rate of 57%.

Many thanks to the SustailNLP program committee for their thorough and thoughtful reviews. We would also like to thank to our panelists and invited speakers whose discussions and talks we strongly believe will make the workshop exciting and memorable.

We are looking forward to the fourth edition of the SustaiNLP workshop!

SustaiNLP Organizers
June 2023

Organizing Committee

Organizers

Nafise Sadat Moosavi, University of Sheffield
Iryna Gurevych, TU Darmstadt
Yufang Hou, IBM Research Ireland
Gyuwan Kim, UC Santa Barbara
Young Jin Kim, Microsoft Research
Tal Schuster, Google Research
Ameeta Agrawal, Portland State University

Program Committee

Program Committee

Hadi Abdi Ghavidel, Bell Canada
Hend Al-khalifa, King Saud University
Ahmed Alajrami, The University of Sheffield
Abhinav Arora, Facebook
Sina Bagheri Nezhad, Portland State University
Mohaddeseh Bastan, Stony Brook University
Samuel Cahyawijaya, HKUST
Roman Grundkiewicz, Microsoft Research
Benjamin Heinzerling, RIKEN AIP and Tohoku University
Amr Hendy, Microsoft
Zhiqi Huang, Tencent AI Lab
Peter Izsak, Intel Labs
Sungho Jeon, Heidelberg Institute for Theoretical Studies
Rishabh Joshi, Google
Mihir Kale, Google
Lis Kanashiro Pereira, Ochanomizu University
Kartik Kannapur, Amazon Web Services
Md. Shahriar Karim, North South University
Constantinos Karouzos, University of Sheffield
Anne Lauscher, University of Hamburg
Ji-ung Lee, UKP, TU Darmstadt
Els Lefever, LT3, Ghent University
Florian Mai, Idiap Research Institute
Jonathan Mamou, Intel Labs
Ekata Mitra, Portland State University
Anhad Mohananey, Google
Yusuke Mori, The University of Tokyo
Marius Mosbach, Saarland University
Sergiu Nisioi, Human Language Technologies Research Center, University of Bucharest
Siddhesh Pawar, Google
Oren Pereg, Emergent AI Lab, Intel Labs
Vikas Raunak, Microsoft
Ines Rehbein, University of Mannheim
Hossein Rouhizadeh, University of Geneva
Edwin Simpson, University of Bristol
Pranaydeep Singh, LT3, University of Ghent
Marek Suppa, Comenius University in Bratislava
Danae Sánchez Villegas, University of Sheffield
Yufei Tao, Portland State University
Lovre Torbarina, doXray
Moshe Wasserblat, Intel Labs
Miles Williams, University of Sheffield
Genta Winata, Bloomberg
Tiezheng Yu, The Hong Kong University of Science and Technology

Table of Contents

KwikBucks: Correlation Clustering with Cheap-Weak and Expensive-Strong Signals
Sandeep Silwal, Sara Ahmadian, Andrew Nystrom, Andrew Mccallum, Deepak Ramachandran and Mehran Kazemi 1

Semantic-Oriented Unlabeled Priming for Large-Scale Language Models
Yanchen Liu, Timo Schick and Hinrich Schtze 32

oBERTa: Improving Sparse Transfer Learning via improved initialization, distillation, and pruning regimes
Daniel Campos, Alexandre Marques, Mark Kurtz and Cheng Xiang Zhai 39

Quick Dense Retrievers Consume KALE: Post Training KullbackLeibler Alignment of Embeddings for Asymmetrical dual encoders
Daniel Campos, Alessandro Magnani and Chengxiang Zhai 59

Lessons on Parameter Sharing across Layers in Transformers
Sho Takase and Shun Kiyono 78

To Asymmetry and Beyond: Structured Pruning of Sequence to Sequence Models for Improved Inference Efficiency
Daniel Campos and Chengxiang Zhai 91

Small is the New Big: Pre-finetuned compact models are better for Asynchronous Active Learning
Dantong Liu, Kaushik Pavani and Sunny Dasgupta 110

ADEPT: Adapter-based Efficient Prompt Tuning Approach for Language Models
Aditya Shah, Surendrabikram Thapa, Aneesh Jain and Lifu Huang 121

NLU on Data Diets: Dynamic Data Subset Selection for NLP Classification Tasks
Jean-michel Attendu and Jean-philippe Corbeil 129

On the Interactions of Structural Constraints and Data Resources for Structured Prediction
Zhisong Zhang, Emma Strubell and Eduard Hovy 147

Can we Pretrain a SotA Legal Language Model on a Budget From Scratch?
Joel Niklaus and Daniele Giofre 158

Is a Video worth n n Images? A Highly Efficient Approach to Transformer-based Video Question Answering
Chenyang Lyu, Tianbo Ji, Yvette Graham and Jennifer Foster 183

How to Unleash the Power of Large Language Models for Few-shot Relation Extraction?
Xin Xu, Yuqi Zhu, Xiaohan Wang and Ningyu Zhang 190

Prompting language models improves performance in imbalanced setting
Jay Mohta 201

KGQA Without Retraining
Nick Mckenna and Priyanka Sen 212

MANER: Mask Augmented Named Entity Recognition for Extreme Low-Resource Languages
Shashank Sonkar, Zichao Wang and Richard Baraniuk 219

Efficient and Interpretable Compressive Text Summarisation with Unsupervised Dual-Agent Reinforcement Learning
Peggy Tang, Junbin Gao, Lei Zhang and Zhiyong Wang 227

Exploring the Effect of Frequency Resolution in FNet
Gregory Szumel, Ghazal Khalighinejad, Rickard Stureborg and Sam Wiseman 239

Towards Adaptable and Interactive Image Captioning with Data Augmentation and Episodic Memory
Aliki Anagnostopoulou, Mareike Hartmann and Daniel Sonntag 245

Corpus Complexity Matters in Pretraining Language Models
Ameeta Agrawal and Suresh Singh 257

PersonaPKT: Building Personalized Dialogue Agents via Parameter-efficient Knowledge Transfer
Xu Han, Bin Guo, Yoon Jung, Benjamin Yao, Yu Zhang, Xiaohu Liu and Chenlei Guo 264

Small Character Models Match Large Word Models for Autocomplete Under Memory Constraints
Ganesh Jawahar, Subhabrata Mukherjee, Debadeepta Dey, Muhammad Abdul-mageed, Laks Lakshmanan, V.s., Caio Mendes, Gustavo De Rosa and Shital Shah 274

Query Encoder Distillation via Embedding Alignment is a Strong Baseline Method to Boost Dense Retriever Online Efficiency
Yuxuan Wang and Lyu Hong 290

Minimalist Entity Disambiguation for Mid-Resource Languages
Benno Kruit 299