NEALT

Northern European Association for Language Technology

NEALT Proceedings Series No. 45

Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)

May 31–2 June, 2021 Reykjavik, Iceland (Online)

Editors: Simon Dobnik and Lilja Øvrelid

NoDaLiDa 2021

23rd Nordic Conference on Computational Linguistics (NoDaLiDa)

Proceedings of the Conference

May 31–2 June, 2021
Reykjavik, Iceland
Online

Front-cover photo of the ongoing volcanic
eruption in Geldingadalir, near Reykjavík
by Kristinn Ingvarsson, University of Iceland.

Published by
Linköping University Electronic Press, Sweden
Linköping Electronic Conference Proceedings, No. 178
NEALT Proceedings Series, No. 45
Indexed in the ACL anthology

ISBN: 978-91-7929-614-8
ISSN: 1650-3686
eISSN: 1650-3740

Sponsors

Message from the General Chair

Welcome to the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa 2021)!

I am a great fan of the NoDaLiDa conference, as a friendly, medium sized conference that offers the opportunity for scientific and social interaction with colleagues from neighbouring countries. When I agreed to serve as the general chair for this years NoDaLiDa I was still relatively optimistic that we could all meet in beautiful Reykjavik, Iceland to enjoy two days of scientific talks, posters and socialising in early June. Unfortunately that turned out to not be possible due to the COVID-19 pandemic. Instead, we are for the first time offering NoDaLiDa as a fully virtual event, free of charge. Even so, I am confident that this years conference will offer the same high-quality program as in previous years and hopefully it can also constitute a meeting place, albeit a digital one, for Northern European NLP researchers in these unusual times.

As in previous editions, the conference features three different types of papers (long, short and demo papers). We received 91 legal submissions, which represents an increase compared to the previous edition of the conference. In total, we accepted 54 papers, which will be presented as 30 oral presentations, 22 posters and 2 demos at the conference. Each paper was reviewed by three experts. We are extremely grateful to the Programme Committee members for their detailed and helpful reviews. Overall, there are 8 oral sessions with talks and two poster sessions organised into themes over the two days, as well as two exciting keynote talks.

I would further like to thank our two great keynote speakers for sharing their work with us: Lucia Specia from Imperial College London will talk about "Disagreement in human evaluation: blame the task not the annotators". Adina Williams from Facebook AI Research (FAIR) will talk about "For Matters Word Order Little MLM". Two exciting talks that complement each other well!

As in previous years, the conference will be preceded by three workshops: Translatology in the Digital Age, NLP for Computer-Assisted Language Learning and Sustainable language representations. I want to thank the workshop organisers for complementing the main program and offering opportunities for in-depth scientific interaction on these diverse and exciting topics.

I would like to thank the entire group of people that made NoDaLiDa 2021 possible. First of all, I would like to thank Beata Megyesi for inviting me to take up this exciting (and at times daunting) role and all her valuable input regarding NEALT and previous editions of NoDaLiDa. I am further indebted to Barbara Plank for her encouragement, for the sharing all the great resources from the last NoDaLiDa and willingly answering questions on all aspects of the conference organisation. I want to thank the program chair committee Jurgita Kapočiūtė-Dzikienė, Mark Fishel, Jón Gudnason, Barbara Plank, Yves Scherrer and Sara Stymne, for working hard on putting the program together. I am particularly grateful to Jurgita Kapočiūtė-Dzikienė, Jón Gudnason, Yves Scherrer and Sara Stymne for their great effort in leading the reviewing process and shepherding papers from submission to a final decision. I could not have done this without you! Special thanks go to the workshop chairs, Hans Moen and Ildikó Pilán, who have done an invaluable job with leading the workshop selection and organisation. A big thanks also to Johannes Bjerva for his work as social media chair and Simon Dobnik for leading the publication efforts that led to this volume, as well as the coordination of the workshop proceedings. Thank you! Finally, my ultimate thanks goes to the local organisation committee and team. Thank you, Hráfn Lóftsson, Anton Karl Ingason and Steinþór Steingrímsson. They are the ones who did all the heavy lifting in the switch to a virtual event and did a truly amazing job!

NoDaLiDa 2021 has received financial support from our generous sponsors, which we would also like to thank here: Lingsoft, Tilde, Mideind and Grammatek. Above all, their support made it possible for us to offer this NoDaLiDa free of charge. I hope that this will open the conference up to an even larger audience of NLP researchers in Northern Europe.

Once again, welcome and I hope you will enjoy the conference!

Lilja Øvrelid

Oslo

May, 2021

Message from the Local Organisers

We were very much looking forward meeting you at the beginning of summer in Reykjavik, Iceland, but due to the COVID-19 pandemic we had to move the conference completely online. This has been a challenge for us, given the fact that NoDaLiDa has never been run online before. We looked at various possible implementations, but at the end we selected a combination of Zoom, YouTube, Gather.town and Trello! Hopefully, we have risen to the challenge and we hope that you will enjoy interesting talks, posters, demos and workshops during theses three days of NoDaLiDa 2021.

Welcome to NoDaLiDa 2021 online!

Organising Committee

General Chair

Lilja Øvrelid, University of Oslo, Norway

Program Committee

Mark Fishel, University of Tartu, Estonia
Jón Guðnason, Reykjavik University, Iceland
Jurgita Kapočiūtė-Dzikienė, Vytautas Magnus University, Lithuania
Barbara Plank, IT University of Copenhagen, Denmark
Yves Scherrer, University of Helsinki, Finland
Sara Stymne, Uppsala University, Sweden

Publication Chair

Simon Dobnik, University of Gothenburg, Sweden

Social Media Chair

Johannes Bjerva, Aalborg University, Copenhagen, Denmark

Workshop Chairs

Hans Moen, University of Turku, Finland
Ildikó Pilán, Norwegian Computing Center, Norway

Workshop Chairs

Hrafn Loftsson, Reykjavik University (chair)
Anton Karl Ingason, University of Iceland (co-chair)
Steinþór Steingrímsson, Arni Magnusson Institute of Icelandic Studies (co-chair)

Invited Speakers

Lucia Specia, Imperial College London
Adina Williams, Facebook AI Research

Reviewers

Yvonne Adesam, University of Gothenburg
Lars Ahrenberg, Linköping University
David Alfter, University of Gothenburg
Tanel Alumäe, Tallinn University of Technology
Krasimir Angelov, University of Gothenburg and Chalmers University of Technology
Rahul Aralikatte, University of Copenhagen
Jeremy Barnes, University of Oslo
Johannes Bjerva, Aalborg University, Copenhagen
Johanna Björklund, Umeå University
Jari Björne, University of Turku
Bernd Bohnet, Google
Marcel Bollmann, University of Copenhagen
Michal Borsky, Reykjavik University
Gerlof Bouma, University of Gothenburg
Johan Boye, KTH Royal Institute of Technology
Maja Buljan, University of Oslo
Marie Candito, Universtité Paris 7 and INRIA
Lin Chen, UIC
Daniel Dakota, Uppsala University
Hercules Dalianis, Stockholm University
Dana Dannells, University of Gothenburg
Miryam de Lhoneux, University of Copenhagen
Simon Dobnik, University of Gothenburg
Adam Ek, University of Gothenburg
Anna Esposito, University of Campania “Luigi Vanvitelli”
Filip Ginter, University of Turku
Ana Gonzalez, University of Copenhagen
Johannes Graën, University of Gothenburg
Hugo Hammer, Oslo Metropolitan University
Mareike Hartmann, University of Copenhagen
Daniel Hershcovich, University of Copenhagen
Nora Hollenstein, ETH Zurich
Bjarte Johansen, Equinor ASA
Richard Johansson, University of Gothenburg
Jenna Kanerva, University of Turku
Jussi Karlgren, Spotify
Andre Kåsen, National Library of Norway
Andreas Søeborg Kirkedal, Interactions LLC
Roman Klinger, University of Stuttgart
Mare Koit, University of Tartu
Dimitrios Kokkinakis, University of Gothenburg
Artur Kulmizev, Uppsala University
Robin Kurtz, KBLab, National Library of Sweden
Andrey Kutuzov, University of Oslo

Veronika Laippala, University of Turku
Krister Lindén, University of Helsinki
Pierre Lison, Norwegian Computing Center
Jan Tore Loenning, University of Oslo
Hrafn Loftsson, Reykjavik University
Juhani Luotolahti, University of Turku
Eydis Magnusdottir, Reykjavik University
Hans Moen, University of Turku
Costanza Navarretta, University of Copenhagen
Anna Nikulásdóttir, Grammatek ehf
Mattias Nilsson, Karolinska Institutet
Joakim Nivre, Uppsala University
Farhad Nooralahzadeh, UZH
Pierre Nugues, Lund University
Emily Öhman, University of Helsinki
Fredrik Olsson, RISE
Robert Östling, Stockholm University
Anthi Papadopoulou, University of Oslo
Victor Petrén Bach Hansen, University of Copenhagen
Eva Pettersson, Uppsala University
Ildiko Pilan, Norwegian Computing Center
Mārcis Pinnis, Tilde
Vinit Ravishankar, University of Oslo
Sagnik Ray Choudhury, University of Copenhagen
Matīss Rikters, University of Tokyo
Fabio Rinaldi, University of Zurich
Eiríkur Rögnvaldsson, University of Iceland
Samuel Rönnqvist, University of Turku
Jack Rueter, University of Helsinki
Magnus Sahlgren, RISE AI
Askars Salimbajevs, Tilde
Marina Santini, RISE
Inguna Skadiņa, Tilde and University of Latvia
Maria Skeppstedt, Stockholm University
Karolina Stanczak, University of Copenhagen
Steinþór Steingrímsson, The Árni Magnússon Institute for Icelandic Studies
Umut Sulubacak, University of Helsinki
Torbjørn Svendsen, Norwegian University of Science and Technology
Gongbo Tang, Uppsala University
Samia Touileb, University of Oslo
Francis M. Tyers, Indiana University Bloomington
Martti Vainio, University of Helsinki
Rob van der Goot, University of Groningen
Daniel Varab, Novo Nordisk and IT University of Copenhagen
Erik Velldal, University of Oslo
Martin Volk, University of Zurich

Elena Volodina, University of Gothenburg
Jürgen Wedekind, University of Copenhagen
Anssi Yli-Jyrä, University of Helsinki
Marcos Zampieri, University of Wolverhampton
Niklas Zechner, University of Gothenburg
Heike Zinsmeister, University of Hamburg

Invited Talks

Lucia Specia: Disagreement in human evaluation: blame the task not the annotators.
It is well known that human evaluators are prone to disagreement and that this is a problem for reliability and reproducibility of evaluation experiments. The reasons for disagreement can fall into two broad categories: (1) human evaluator, including under-trained, under-incentivised, lacking expertise, or ill-intended individuals, e.g., cheaters; and (2) task, including ill-definition, poor guidelines, sub-optimal setup, or inherent complexity or subjectivity. While in an ideal evaluation experiment many of these elements will be controlled for, in this talk I will argue that task complexity and subjectivity are much harder issues and that in some cases agreement cannot and should not be expected. I will cover several evaluation experiments on tasks with variable degrees of complexity and subjectivity, discuss their levels of disagreement along with other issues. I hope this will lead to an open discussion on possible strategies and directions to address this problem.

Adina Williams: For Matters Word Order Little MLM.
One possible explanation for the impressive performance of masked language models (MLMs) is that they can learn to represent the syntactic structures prevalent in classical NLP pipelines. Were this correct, we would expect that fine-tuning such models on tasks requiring syntactic structure would lead them to be sensitive to word order at inference time. To address this question, we permute example word order at several steps in the pipeline—during fine-tuning, evaluation, and/or pre-training—and measure the results. We find that permuting word order during fine-tuning has remarkably little effect on downstream performance for several purportedly syntax sensitive NLU tasks (including NLI). Next, we pre-train MLMs on examples with randomly shuffled word order, and find that these models still achieve high accuracy (even after unpermuted fine-tuning) on many downstream tasks—including tasks specifically designed to be challenging for models that ignore word order. Our results show that the success of MLM pre-training is largely due to distributional information not any knowledge of word order per se, and underscores the importance of curating challenging evaluation datasets that require deeper syntactic knowledge.