NEALT

Northern European Association for Language Technology

NEALT Proceedings Series No. 57

Proceedings of the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025)

March 3 – 4, 2025

Tallinn, Estonia

Editors: Richard Johansson and Sara Stymne

NoDaLiDa/Baltic-HLT 2025

Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025)

Proceedings of the Conference

March 3-4, 2025

The NoDaLiDa/Baltic-HLT 2025 organizers gratefully acknowledge the support from the following sponsors.

©2025 University of Tartu Library

Front-cover photo: Rasmus Jurkatam

Published by:

University of Tartu Library, Estonia
NEALT Proceedings Series, No. 57
Indexed in the ACL Anthology

ISBN: 978-9908-53-109-0
ISSN: 1736-8197 (Print)
ISSN: 1736-6305 (Online)

Volume Editors:
Richard Johansson and Sara Stymne

Message from the General Chair

Welcome to the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025) to be held in beautiful Tallinn, Estonia, on March 2–5, 2025.

It is 48 years since the first NoDaLiDa was held and 21 years since the first Baltic HLT was held. Now, for the first time, the two major conferences on computational linguistics and language technology in the Nordic and Baltic regions have joined forces as a joint event. Both conferences aim to bring together researchers in the Nordic and Baltic countries interested in any aspect related to human language and speech technology. As a joint event, we extended the conference with one extra workshop day so that we have one day before and one day after the main two-day conference. It is a great honor for me to serve as the general chair of this joint event.

We solicited three different types of papers (long, short, and demo papers) and received 127 valid submissions, of which 4 were withdrawn during the process. In total, we accepted 81 papers (acceptance rate: 66%; long papers 65%, short papers: 66%, demos 80%), which will be presented as 43 oral presentations, 34 posters, and 4 demos. More than half of the accepted papers are student papers, in which the first author is a student (29 long, 19 short, and 2 demo papers). Each paper was reviewed by three experts. We are extremely grateful to the 155 Programme Committee members for their detailed and helpful reviews.

The 81 accepted papers are organized into 12 oral sessions and 2 poster and demo sessions. In addition to these regular sessions, the conference program includes three keynote talks. We would like to extend our gratitude to the keynote speakers for agreeing to present their work at NoDaLiDa/Baltic-HLT. Arianna Bisazza from the University of Groningen will talk on the topic of "Not all Language Models Need to be Large: Studying Language Evolution and Acquisition with Modern Neural Networks." Dirk Hovy from Bocconi University will talk about "The Illusion of Understanding – Unpacking the True Capabilities of Language Models." Arvi Tavast from the Institute of the Estonian Language will talk about "No Sex, No Future: On the Status of Estonian in a Changing World," continuing the NoDaLiDa tradition of featuring a presentation about the local language.

The main conference is complemented by 6 workshops on a diverse set of topics. On March 2, preceding the main conference: Resources and representations for under-resourced languages and domains (RESOURCEFUL-2025); Nordic-Baltic Responsible Evaluation and Alignment of Language models (NB-REAL); and The 1st Workshop on Ecology, Environment, and Natural Language Processing (NLP4Ecology). On March 5, after the main conference: Constraint Grammar and Finite State NLP – Rule-based and hybrid methods and tools for user communities; The 13th Workshop on Natural Language Processing for Computer Assisted Language Learning (NLP4CALL); and Automatic Assessment of Atypical Speech (AAAS). The workshop themes illustrate the breadth of topics that can be found in language technology, and we are extremely happy and grateful to the workshop organizers for complementing the main program.

I would like to thank the entire team that made NoDaLiDa/Baltic-HLT possible. I was honored to receive the invitation to serve as the general chair from Jörg Tiedemann and the NEALT board; thank you for trusting me in this role. My deepest gratitude goes to the organizing committee. Thank you to the program chair committee Daniel Hershcovich, Jenna Kanerva, Pierre Lison, and Andrius Utka, for working hard on putting the program together, especially for your great effort in leading the reviewing process and shepherding papers from submission to the final decision. Thank you to the program chair advisors Mark Fišel and Inguna Skadiņa, for your valuable advice about previous editions of NoDaLiDa and Baltic-HLT. Thank you to Richard Johansson for leading the publication efforts that led to this volume, as well as the coordination of the workshop proceedings. Thank you to the workshop chairs,

Normunds Grūzītis and Samia Touileb, for leading the workshop selection. Thank you to Mike Zhang, our social media chair, for all your posts and for spreading information about the conference. My ultimate thank you goes to the local organizer team, Helen Kaljumäe, Merily Remma, and Kadri Vare, for a truly amazing job; the conference wouldn't have happened without your effort! It was a pleasure to work together.

On behalf of the organizing committee, we would like to thank the NoDaLiDa/Baltic-HLT sponsors for their generous financial support that helped us organize an affordable conference. We would also like to thank all the conference speakers and participants. Your interactions and enthusiasm are what will make the actual conference into a forum for fruitful conversations and discussions that contribute to connections for years to come.

Welcome, and I hope you enjoy the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies

Sara Stymne
Uppsala
March 2025

Message from the Local Organizers

When we began planning this event, we initially set out to host Baltic-HLT, a key conference for the Baltic language technology community. However, it ended up that this year marks the first time that Baltic-HLT and NoDaLiDa have been brought together, uniting two established traditions into one joint conference. In an era shaped by rapid advancements in language technologies – especially the rise of large language models – collaboration across regions and disciplines is more important than ever. Together we can address common challenges and ensure that every language, large or small, has a place in the digital future. We are pleased to welcome you to Tallinn, where cultural richness and innovation come together! We hope your time here brings new ideas, valuable connections, and leaves you with great memories.

Thank you for everyone for being part of NoDaLiDa/Baltic-HLT 2025, and we wish you a wonderful conference!

Organizing Committee

General Chair

Sara Stymne, Uppsala University, Sweden

Program Chair Advisors

Mark Fišel, University of Tartu, Estonia
Inguna Skadiņa, University of Latvia/Tilde, Latvia

Program Chairs

Daniel Hershcovich, University of Copenhagen, Denmark
Jenna Kanerva, University of Turku, Finland
Pierre Lison, Norwegian Computing Centre, Norway
Andrius Utka, Vytautas Magnus University, Lithuania

Workshop Chairs

Normunds Grūzītis, University of Latvia, Latvia
Samia Touileb, University of Bergen, Norway

Publication Chair

Richard Johansson, Chalmers University of Technology and University of Gothenburg, Sweden

Social Media Chair

Mike Zhang, Aalborg University, Denmark

Local Organizers

Helen Kaljumäe, Institute of the Estonian Language, Estonia
Merily Remma, Institute of the Estonian Language, Estonia
Kadri Vare, Institute of the Estonian Language, Estonia

Reviewers

Eleri Aedmaa, Manex Agirrezabal, Lars Ahrenberg, David Alfter, Ali Al-Laith, Sven Aller, Kais Allkivi, Tanel Alumäe, Mark Anderson, Mattias Appelgren, Ilze Auzina, Eduard Barbu, Guntis Barzdins, Ali Basirat, Elisa Bassignana, Timo Baumann, Aleksandrs Berdicevskis, Johanna Björklund, Jari Björne, Gerlof Bouma, Johan Boye, Stephanie Brandl, Chloé Braud, Micaella Bruton, Maja Buljan, Laura Cabello, Lucas Georges Gabriel Charpentier, Yiyi Chen, Mathias Creutz, David Dale, Vera Danilova, Iben Nyholm Debess, Hannah Devinney, Ruchira Dhar, Stefanie Dipper, Simon Dobnik, Lucia Donatelli, Aleksei Dorkin, Luise Dürlich, Jens Edlund, Antske Fokkens, Filip Ginter, Evangelia Gogoulou, Rob van der Goot, Tamás Grósz, Hugo Lewi Hammer, Stefan Heinrich, Aron Henriksson, Erik Henriksson, Oskar Holmström, David House, Christine Howes, Guimin Hu, Maciej Janicki, Moa Johansson, Arne Jönsson, Heiki-Jaan Kaalep, Emil Kalbaliyev, Danguolė Kalinauskaitė, Kaarel Kaljurand, Amanda Kann, Antti Kanner, Jurgita Kapočiūtė-Dzikienė, Jussi Karlgren, Andreas Kirkedal, Mare Koit, Tomas Krilavičius, Marco Kuhlmann, Jenny Kunz, Murathan Kurfalı, Mikko Kurimo, Robin Kurtz, Andrey Kutuzov, Hele-Andra Kuulmets, Harm Lameris, Ekaterina Lapshinova-Koltunski, Heather Lent, Krister Lindén, Ellinor Lindqvist, Sharid Loáiciga, Agnes Luhtaru, Petter Mæhlum, Giacomo Magnifico, Arianna Masciolini, Timothee Mickus, Vladislav Mikhailov, Kirill Milintsevich, Hans Moen, Anssi Moisio, Kadri Muischnek, Kaili Müürisep, Arild Brandrud Næss, Costanza Navarretta, Anna Björk Nikulasdottir, Joakim Nivre, Bill Noble, Emily Öhman, Siim Orasmaa, Jim O'Regan, Petya Osenova, Robert Östling, Siddhesh Milind Pawar, Bolette S. Pedersen, Qiwei Peng, Eva Pettersson, Mārcis Pinnis, Flammie A Pirinen, Taido Purason, Alessandro Raganato, Liisa Rätsep, Sebastian Reimann, Matīss Rikters, Egil Rønningstad, Ahmed Ruby, Askars Salimbajevs, David Samuel, Ricardo Muñoz Sánchez, David Sasu, Baiba Valkovska Saulīte, Denitsa Saynova, Barbara Scalvini, Yves Scherrer, Djamé Seddah, Kiril Ivanov Simov, Kairit Sirts, Maria Skeppstedt, Somnath, Steinþór Steingrímsson, Felix Stollenwerk, Antti Suni, Torbjørn Svendsen, Monorama Swain, Nina Tahmasebi, Gongbo Tang, Timothy R Tangherlini, Arvi Tavast, Jörg Tiedemann, Crina Tudor, Dennis Ulmer, Teemu Vahtola, Jurgita Vaičenonienė, Thomas Vakili, Daniel Varab, Erik Velldal, Martin Volk, Elena Volodina, Fredrik Wahlberg, Nicholas Thomas Walker, Sondre Wold, Lisa Yankovskaya, Huiling You, Oreen Yousuf, Niklas Zechner, Ingo Ziegler, Heike Zinsmeister

Invited Talk: Not all Language Models Need to be Large – Studying Language Evolution and Acquisition with Modern Neural Networks

Arianna Bisazza
University of Groningen

Why do languages look the way they do? And what makes us so good at learning language as we grow up? Since the early days of connectionism, outstanding questions about human language have been investigated by means of simulations involving small neural networks (NNs) and toy languages. Is this still possible and meaningful in the age of Large pre-trained Language Models (LLMs)?

In this talk, I'll propose that modern NNs can indeed be a valuable tool to simulate and study processes of language evolution and acquisition. This, however, requires having control of training data, model architecture, and learning setup, which is typically not possible with LLMs.

I will then present two lines of research following these principles, namely: (1) simulating language change using small NN-agents that learn to communicate with pre-defined artificial languages, and (2) simulating the acquisition of syntax by training LMs on child-directed language. I'll end with a discussion of the value of interdisciplinarity and the importance of experimenting in controlled setups, rather than focusing all our research efforts on the evaluation of LLMs.

Invited Talk: The Illusion of Understanding – Unpacking the True Capabilities of Language Models

Dirk Hovy
Bocconi University

The rapid development of large language models in recent years has transformed the field of NLP. Many people are concerned that it has trivialized the field or even rendered it obsolete. In this talk, I'll argue that neither is true: NLP has a long way to go, and LLMs are the most recent in a long line of methods that have advanced the field. LLMs have freed us from many of the nitty-gritty details that previously hampered NLP research, allowing us to focus on larger and more interesting questions.

One of the most fundamental questions is what it means to "understand" language. In a world where AI can generate anything from translations to poetry and code, it's easy to believe these models genuinely understand us. However, despite its linguistic abilities, today's generative AI still resembles a skilled mimic rather than a genuine linguist. We will look at thought experiments and real-world examples to demonstrate the limitations of statistical models' knowledge, their inability to grasp context and nuance, and the dangers of overestimating their abilities. I will emphasize the theoretical and practical implications for future language technology, with a focus on social context. Drawing on philosophy, linguistics, and NLP history, we will investigate what it truly means to 'understand' a language beyond the words and the implications for safety and utility in LLMs.

Invited Talk: No Sex, no Future – On the Status of Estonian in a Changing World

Arvi Tavast
Institute of the Estonian Language

Apart from well-known anecdotes about the absence of gender marking and future tense, the most peculiar feature of Estonian is its number of speakers. Being one of the smallest fully functional languages in the world, it is a source of pride for its speakers, as well as a central part of their identity. The resulting puristic attitudes towards language also enjoy strong legal support. One of the enablers of this ideological stance is the channel metaphor of communication: that language as a system exists independently of its speakers, and communication works in virtue of using a shared code to encode and decode messages. This metaphor is still going strong in folk linguistics despite all evidence to the contrary, including recent advances in language modelling. A completely different reading for the title of the talk is provided by more recent learning- and prediction-based accounts of why we understand each other. Language, like any naturally evolving system, is vitally dependent on the random variability that is so conveniently present in linguistic data. This makes openness to new information a precondition to having a future, also for languages.

Table of Contents

Annotating and Classifying Direct Speech in Historical Danish and Norwegian Literary Texts
Ali Al-Laith, Alexander Conroy, Kirstine Nielsen Degn, Jens Bjerring-Hansen and Daniel Hershcovich . . . 1

Diachronic Analysis of Phrasal Verbs in English Scientific Writing
Diego Alves . . . 8

Applying and Optimising a Multi-Scale Probit Model for Cross-Source Text Complexity Classification and Ranking in Swedish
Elsa Andersson, Johan Falkenjack and Arne Jönsson . . . 17

Playing by the Rules: A Benchmark Set for Standardized Icelandic Orthography
Bjarki Ármannsson, Hinrik Hafsteinsson, Jóhannes B. Sigtryggsson, Atli Jasonarson, Einar Freyr Sigurðsson and Steinþór Steingrímsson . . . 28

An Icelandic Linguistic Benchmark for Large Language Models
Bjarki Ármannsson, Finnur Ágúst Ingimundarson and Einar Freyr Sigurðsson . . . 37

Transfer-Learning German Metaphors Inspired by Second Language Acquisition
Maria Berger . . . 48

Comparative Concepts or Descriptive Categories: a UD Case study
Matthieu Pierre Boyer and Mathieu Dehouck . . . 55

Investigating the effectiveness of Data Augmentation and Contrastive Learning for Named Entity Recognition
Noel Chia, Ines Rehbein and Simone Paolo Ponzetto . . . 66

Comparing Human and Machine Translations of Generative Language Model Evaluation Datasets
Sander Bijl de Vroe, George Stampoulidis, Kai Hakala, Aku Rouhe, Mark van Heeswijk and Jussi Karlgren . . . 80

GliLem: Leveraging GliNER for Contextualized Lemmatization in Estonian
Aleksei Dorkin and Kairit Sirts . . . 86

Comparative analysis of optical character recognition methods for Sámi texts from the National Library of Norway
Tita Enstad, Trond Trosterud, Marie Iversdatter Røsok, Yngvil Beyer and Marie Roald . . . 98

LAG-MMLU: Benchmarking Frontier LLM Understanding in Latvian and Giriama
Naome A. Etori, Arturs Kanepajs, Kevin Lu and Randu Karisa . . . 109

Better Benchmarking LLMs for Zero-Shot Dependency Parsing
Ana Ezquerro, Carlos Gómez-Rodríguez and David Vilares . . . 121

Optimizing Estonian TV Subtitles with Semi-supervised Learning and LLMs
Artem Fedorchenko and Tanel Alumäe . . . 136

Modeling Multilayered Complexity in Literary Texts
Pascale Feldkamp, Márton Kardos, Kristoffer Nielbo and Yuri Bizzoni . . . 142

Does Preprocessing Matter? An Analysis of Acoustic Feature Importance in Deep Learning for Dialect Classification
Lea Fischbach, Caroline Kleen, Lucie Flek and Alfred Lameli . . . 159

Language of the Swedish Manosphere with Swedish FrameNet
Emilie Marie Carreau Francis 170

Hotter and Colder: A New Approach to Annotating Sentiment, Emotions, and Bias in Icelandic Blog Comments
Steinunn Rut Friðriksdóttir, Dan Saattrup Nielsen and Hafsteinn Einarsson 181

Towards large-scale speech foundation models for a low-resource minority language
Yaroslav Getman, Tamás Grósz, Katri Hiovain-Asikainen, Tommi Lehtonen and Mikko Kurimo 192

OpusDistillery: A Configurable End-to-End Pipeline for Systematic Multilingual Distillation of Open NMT Models
Ona de Gibert, Tommi Nieminen, Yves Scherrer and Jörg Tiedemann 201

Mind the Gap: Diverse NMT Models for Resource-Constrained Environments
Ona de Gibert, Dayyán O'Brien, Dušan Variš and Jörg Tiedemann 209

Testing relevant linguistic features in automatic CEFR skill level classification for Icelandic
Isidora Glišić, Caitlin Laura Richter and Anton Karl Ingason 217

MorSeD: Morphological Segmentation of Danish and its Effect on Language Modeling
Rob van der Goot, Anette Jensen, Emil Allerslev Schledermann, Mikkel Wildner Kildeberg, Nicolaj Larsen, Mike Zhang and Elisa Bassignana 223

Opinion Units: Concise and Contextualized Representations for Aspect-Based Sentiment Analysis
Emil Häglund and Johanna Björklund 230

Aligning Language Models for Icelandic Legal Text Summarization
Þórir Hrafn Harðarson, Hrafn Loftsson and Stefán Ólafsson 241

Question-parsing with Abstract Meaning Representation enhanced by adding small datasets
Johannes Heinecke, Maria Boritchev and Frédéric Herledan 252

FinerWeb-10BT: Refining Web Data with LLM-Based Line-Level Filtering
Erik Henriksson, Otto Tarkka and Filip Ginter 258

Margins in Contrastive Learning: Evaluating Multi-task Retrieval for Sentence Embeddings
Tollef Emil Jørgensen and Jens Breitung 269

Database of Latvian Morphemes and Derivational Models: ideas and expected results
Andra Kalnača, Tatjana Pakalne and Kristīne Levāne-Petrova 279

Localizing AI: Evaluating Open-Weight Language Models for Languages of Baltic States
Jurgita Kapočiūtė-Dzikienė, Toms Bergmanis and Mārcis Pinnis 287

How Aunt-Like Are You? Exploring Gender Bias in the Genderless Estonian Language: A Case Study
Elisabeth Kaukonen, Ahmed Sabir and Rajesh Sharma 296

Estonian isolated-word text-to-speech synthesiser
Indrek Kiissel, Liisi Piits, Heete Sahkai, Indrek Hein, Liis Ermus and Meelis Mihkla 302

BiaSWE: An Expert Annotated Dataset for Misogyny Detection in Swedish
Kätriin Kukk, Danila Petrelli, Judit Casademont, Eric J. W. Orlowski, Michal Dzielinski and Maria Jacobson 307

Predictability of Microsyntactic Units across Slavic Languages: A translation-based Study
Maria Kunilovskaya, Iuliia Zaitova, Wei Xue, Irina Stenger and Tania Avgustinova 313

Train More Parameters But Mind Their Placement: Insights into Language Adaptation with PEFT
Jenny Kunz 323

SweSAT-1.0: The Swedish University Entrance Exam as a Benchmark for Large Language Models
Murathan Kurfalı, Shorouq Zahra, Evangelia Gogoulou, Luise Dürlich, Fredrik Carlsson and Joakim Nivre 331

How Well do LLMs know Finno-Ugric Languages? A Systematic Assessment
Hele-Andra Kuulmets, Taido Purason and Mark Fishel 340

Mapping Faroese in the Multilingual Representation Space: Insights for ASR Model Optimization
Dávid í Lág, Barbara Scalvini and Jon Gudnason 354

Towards a Derivational Semantics Resource for Latvian
Ilze Lokmane, Mikus Grasmanis, Agute Klints, Gunta Nešpore-Bērzkalne, Pēteris Paikens, Lauma Pretkalniņa, Laura Rituma, Madara Stāde and Evelīna Tauriņa 359

Poro 34B and the Blessing of Multilinguality
Risto Luukkonen, Jonathan Burdge, Elaine Zosa, Aarne Talman, Ville Komulainen, Väinö Hatanpää, Peter Sarlin and Sampo Pyysalo 367

Can summarization approximate simplification? A gold standard comparison
Giacomo Magnifico and Eduard Barbu 383

A Comparative Study of PEFT Methods for Python Code Generation
Johanna Männistö, Joseph Attieh and Jörg Tiedemann 390

A Collection of Question Answering Datasets for Norwegian
Vladislav Mikhailov, Petter Mæhlum, Victoria Ovedie Chruickshank Langø, Erik Velldal and Lilja Øvrelid 397

Incorporating Target Fuzzy Matches into Neural Fuzzy Repair
Tommi Nieminen, Jörg Tiedemann and Sami Virpioja 408

Constructions and Strategies in Universal Dependencies
Joakim Nivre 419

Finnish SQuAD: A Simple Approach to Machine Translation of Span Annotations
Emil Nuutinen, Iiro Rastas and Filip Ginter 424

How to Tune a Multilingual Encoder Model for Germanic Languages: A Study of PEFT, Full Fine-Tuning, and Language Adapters
Romina Oji and Jenny Kunz 433

Match 'em: Multi-Tiered Alignment for Error Analysis in ASR
Phoebe Parsons, Knut Kvale, Torbjørn Svendsen and Giampiero Salvi 440

Adding Metadata to Existing Parliamentary Speech Corpus
Phoebe Parsons, Per Erik Solberg, Knut Kvale, Torbjørn Svendsen and Giampiero Salvi 448

Paragraph-Level Machine Translation for Low-Resource Finno-Ugric Languages
Dmytro Pashchenko, Lisa Yankovskaya and Mark Fishel 458

Evaluating LLM-Generated Explanations of Metaphors – A Culture-Sensitive Study of Danish
Bolette S. Pedersen, Nathalie Sørensen, Sanni Nimb, Dorte Haltrup Hansen, Sussi Olsen and Ali Al-Laith 470

Tokenization on Trial: The Case of Kalaallisut–Danish Legal Machine Translation
Esther Ploeger, Paola Saucedo, Johannes Bjerva, Ross Deans Kristensen-McLachlan and Heather Lent 480

The Roles of English in Evaluating Multilingual Language Models
Wessel Poelman and Miryam de Lhoneux 492

Revisiting Projection-based Data Transfer for Cross-Lingual Named Entity Recognition in Low-Resource Languages
Andrei Politov, Oleh Shkalikov, Rene Jäkel and Michael Färber 499

Empathy vs Neutrality: Designing and Evaluating a Natural Chatbot for the Healthcare Domain
Cristina Reguera-Gómez, Denis Paperno and Maaike H. T. de Boer 508

Assessed and Annotated Vowel Lengths in Spoken Icelandic Sentences for L1 and L2 Speakers: A Resource for Pronunciation Training
Caitlin Laura Richter, Kolbrún Friðriksdóttir, Kormákur Logi Bergsson, Erik Anders Maher, Ragnheiður María Benediktsdóttir and Jon Gudnason 518

The BRAGE Benchmark: Evaluating Zero-shot Learning Capabilities of Large Language Models for Norwegian Customer Service Dialogues
Mike Riess and Tollef Emil Jørgensen 525

Mixed Feelings: Cross-Domain Sentiment Classification of Patient Feedback
Egil Rønningstad, Lilja Charlotte Storset, Petter Mæhlum, Lilja Øvrelid and Erik Velldal 537

The Impact of Copyrighted Material on Large Language Models: A Norwegian Perspective
Javier de la Rosa, Vladislav Mikhailov, Lemei Zhang, Freddy Wetjen, David Samuel, Peng Liu, Rolv-Arild Braaten, Petter Mæhlum, Magnus Breder Birkenes, Andrey Kutuzov, Tita Enstad, Hans Christian Farsethås, Svein Arne Brygfjeld, Jon Atle Gulla, Stephan Oepen, Erik Velldal, Wilfred Østgulen, Lilja Øvrelid and Aslak Sira Myhre 544

Encoder vs Decoder: Comparative Analysis of Encoder and Decoder Language Models on Multilingual NLU Tasks
Dan Saattrup Nielsen, Kenneth Enevoldsen and Peter Schneider-Kamp 561

Small Languages, Big Models: A Study of Continual Training on Languages of Norway
David Samuel, Vladislav Mikhailov, Erik Velldal, Lilja Øvrelid, Lucas Georges Gabriel Charpentier, Andrey Kutuzov and Stephan Oepen 573

Rethinking Low-Resource MT: The Surprising Effectiveness of Fine-Tuned Multilingual Models in the LLM Age
Barbara Scalvini, Iben Nyholm Debess, Annika Simonsen and Hafsteinn Einarsson 609

Prompt Engineering Enhances Faroese MT, but Only Humans Can Tell
Barbara Scalvini, Annika Simonsen, Iben Nyholm Debess and Hafsteinn Einarsson 622

Interactive maps for corpus-based dialectology
Yves Scherrer and Olli Kuparinen 634

Profiling Bias in LLMs: Stereotype Dimensions in Contextual Word Embeddings
Carolin M. Schuster, Maria-Alexandra Roman, Shashwat Ghatiwala and Georg Groh 639

Entailment Progressions: A Robust Approach to Evaluating Reasoning Within Larger Discourse
Rishabh Shastry, Patricia Chiril, Joshua Charney and David Uminsky 651

Generative AI for Technical Writing: Comparing Human and LLM Assessments of Generated Content
Karen de Souza, Alexandre Nikolaev and Maarit Koponen 661

MC-19: A Corpus of 19th Century Icelandic Texts
Steinþór Steingrímsson, Einar Freyr Sigurðsson and Atli Jasonarson 680

Surface-Level Morphological Segmentation of Low-resource Inuktitut Using Pre-trained Large Language Models
Mathias Stenlund, Hemanadhan Myneni and Morris Riedel 688

The Devil's in the Details: the Detailedness of Classes Influences Personal Information Detection and Labeling
Maria Irena Szawerna, Simon Dobnik, Ricardo Muñoz Sánchez and Elena Volodina 697

Braxen 1.0
Christina Tånnander and Jens Edlund 709

Temporal Relation Classification: An XAI Perspective
Sofia Elena Terenziani 714

Benchmarking Abstractive Summarisation: A Dataset of Human-authored Summaries of Norwegian News Articles
Samia Touileb, Vladislav Mikhailov, Marie Ingeborg Kroka, Lilja Øvrelid and Erik Velldal . . 729

Efficient Elicitation of Fictitious Nursing Notes from Volunteer Healthcare Professionals
Jesper Vaaben Bornerup and Christian Hardmeier 739

Analyzing the Effect of Linguistic Instructions on Paraphrase Generation
Teemu Vahtola, Songbo Hu, Mathias Creutz, Ivan Vulić, Anna Korhonen and Jörg Tiedemann 755

SweClinEval: A Benchmark for Swedish Clinical Natural Language Processing
Thomas Vakili, Martin Hansson and Aron Henriksson 767

Dialectal treebanks and their relation with the standard variety: The case of East Cretan and Standard Modern Greek.
Socrates Vakirtzian, Vivian Stamou, Yannis Kazos and Stella Markantonatou 776

Danoliteracy of Generative Large Language Models
Søren Vejlgaard Holm, Lars Kai Hansen and Martin Carsten Nielsen 785

NorEventGen: generative event extraction from Norwegian news
Huiling You, Samia Touileb, Erik Velldal and Lilja Øvrelid 801

SnakModel: Lessons Learned from Training an Open Danish Large Language Model
Mike Zhang, Max Müller-Eberstein, Elisa Bassignana and Rob van der Goot 812

Got Compute, but No Data: Lessons From Post-training a Finnish LLM
Elaine Zosa, Ville Komulainen and Sampo Pyysalo 826