

SwissText 2026

**Proceedings of the 11th Edition of the
Swiss Text Analytics Conference**

June 10, 2026
University of Zurich
Zurich, Switzerland

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


Sponsors



swissdox.ch



Partners



Datalab - The ZHAW Data Science Laboratory



Fachhochschule Graubünden



swiss made software



©2026 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

Introduction

This year, the 11th SwissText 2026 was hosted at the University of Zurich (UZH) on 10th June in Zurich, Oerlikon, and co-organized by SwissNLP, with over 100 participants attending keynotes, talks, and the poster session.

Our call for papers asked for contributions in four major tracks: Scientific Track (archival), Corpus Track (archival), Applied Track (non-archival), and Demonstration Track (non-archival). We received a total of 42 submissions, and selected 12 for oral presentations, and 25 as posters.

The SwissText 2026 proceedings contain:

- 20 papers for the **Scientific Track**, with technical research papers from the international scientific community, including corpus- and benchmark-related research papers with a focus on Swiss languages from the scientific community and industry,
- 2 papers for the **Corpus Track**, with Swiss-related NLP datasets.

We would like to thank our keynote speakers Alexandra Birch and Valentina Pyatkin. Their perspectives and contributions were much appreciated.

We are grateful to our sponsors and partners, who supported us. In addition, we would like to express our sincere gratitude to everyone who has contributed to this conference. Thanks to conference fellow members of the organizing committee, and the program committee for their excellent work.

Yours sincerely,

Rico Sennrich, Gerold Schneider, Yingqiang Gao, Tilia Ellendorff
and Jannis Vamvas, University of Zurich

Mark Cieliebak, Zurich University of Applied Sciences

General Chair:

Rico Sennrich, University of Zurich

Program Chairs:

Gerold Schneider, University of Zurich

Yingqiang Gao, University of Zurich

Tilia Ellendorff, University of Zurich

Jannis Vamvas, University of Zurich

Local Organizers:

Jannis Vamvas, University of Zurich

Yingqiang Gao, University of Zurich

Michelle Wastl, University of Zurich

Gerold Schneider, University of Zurich

Tilia Ellendorff, University of Zurich

Maria Christina Panagiotopoulou, University of Zurich

Program Committee:

Elliott Ash, Fernando Benites, Miriam Butt, Mark Cieliebak, Simon Clematide, Yingqiang Gao, Jonathan Gerber, Brigitte Krenn, Wolfgang Menzel, Margot Mieskes, Andrei Popescu-Belis, Annette Rios, Josef Ruppenhofer, Tanja Samardzic, Tatjana Scheffler, Helmut Schmid, Gerold Schneider, Manfred Stede, Don Tuggener, Ahmet Yavuz Uluslu, Junling Wang, Egon Werlen, Torsten Zesch, Heike Zinsmeister, Vilém Zouhar

Invited Speakers:

Alexandra Birch, School of Informatics, University of Edinburgh / Aveni.ai

Valentina Pyatkin, Allen Institute for AI / ETH AI Center

Volunteers:

Zachary Hopton

Annalena Alber

Sophia Conrad

Kirill Semenov

Jana Massoud

Keynote Talk
**From Metrics to Models: Reproducible Research in
Multilingual NLP**

Alexandra Birch

School of Informatics, University of Edinburgh / Aveni.ai

Abstract: *Reproducibility in multilingual NLP depends not only on open models, but also on evaluation methods that produce reliable scientific conclusions. In this talk, I connect insights from Extrinsic Evaluation of MT Metrics, which reframed metric assessment around downstream utility and robustness, with EuroLLM, an effort to make large-scale multilingual model development more transparent and accessible. Together, these works highlight two complementary pillars of reproducible research: building systems that others can replicate and measuring them in ways that others can trust and verify. I will discuss how this perspective has shaped the field's move toward open multilingual models, stronger evaluation practices, and more reproducible progress in AI.*

Keynote Talk
**Lessons from Building, Training and Evaluating Open
Language Models**

Valentina Pyatkin
Allen Institute for AI / ETH AI Center

Abstract: *This talk will outline the current landscape of fully open, open-weight, and closed language models, examining what distinguishes fully open models like Apertus from open-weights releases, and discussing the research and engineering considerations that shape open LLM development. Using OLMo and Apertus as case studies, I will walk through the full stack of decisions involved in training modern LLMs: from team dynamics and data curation to compute requirements and post-training methodology. The talk concludes with a discussion of the current capabilities and limitations of open models, with a particular focus on instruction-following capabilities as measured by IFBench.*

Table of Contents

1 Scientific Track	1
<i>Reinforcement Learning for Latent-Space Thinking in LLMs</i> Enes Özeren and Matthias Assenmacher	2
<i>Enhancing Retrieval via Cognitively Motivated Document Expansion</i> Giacomo Loss, Andreas Stephan and Matthias Assenmacher	18
<i>An Efficient Approach for Answering Not Readily Attainable Questions for RAG-based Applications</i> Zhengdao Chen, Christian Heumann and Matthias Assenmacher	30
<i>Concept Extraction and Webb’s Depth of Knowledge: Comparing LLM Question Generation Pipelines for Educational Assessment</i> Dmitriy An, Andrew Paice, Petra Müller-Csernetzky and Aliaksei Andrushevich	53
<i>Optimizing Large Language Models for Robust Domain-Specific Text-to-SQL: From Prompting to Preference Alignment</i> Noah Hampp, Katya Mirylenka and Michael Glass	64
<i>Call Support Copilot: A Reproducible Multimodal System for Speech Emotion Recognition, Intent Understanding, and Agent Assistance</i> Khoshimov Rakhmatillokhon, Dmitry Rudshin and Yanyang Luo	76
<i>Text vs. Phoneme Intermediates for Low-Resource Swiss German</i> Reza Kakooee, Vincenzo Timmel, Daniel Olivier Perruchoud, Michael Graber and Manfred Vogel	83
<i>Extending the Contact Hypothesis: Cross-Linguistic Evaluation of Religion and Nationality Bias When Prompting LLMs in German and Icelandic</i> Ikae Catherine, Guðbjörg Linda Rafnsdóttir, Ragna Kemp Haraldsdóttir and Mascha Kurpicz-Briki	91
<i>Robust Language Identification for Romansh Varieties</i> Charlotte Model, Sina Ahmadi and Jannis Vamvas	102
<i>Skill Extraction from Resumes and Job Offers across Six Languages</i> Laura Vásquez-Rodríguez, Bertrand Audrin, Samuel Michel, Samuele Galli, Julneth Rogenhofer, Jacopo Negro Cusa and Lonneke van der Plas	112
<i>RUMLEM: A Dictionary-Based Lemmatizer for Romansh</i> Dominic P. Fischer, Zachary Hopton and Jannis Vamvas	126
<i>Graph-Augmented LLMs for Swiss MP Ideology Prediction</i> Yifei Yuan, Luis Salamanca, Sophia Schlosser and Laurence Brandenberger	134
<i>Extracting Article-Level Legal Dependencies from Swiss Federal Law using LLMs</i> Steven Cho, Anna Kiepora, Jessica Lam and Richard Hahnloser	147
<i>Data Augmentation for Historical NER: A Systematic Comparison of Lexical and LLM-based Approaches</i> Léa Blinière, Maud Ehrmann, Emanuela Boros, Simon Clematide and Frederic Kaplan	155
<i>Code-Switching Detection in Multilingual Child Speech with SwissBERT</i> Agnese D’Angelo, Sina Ahmadi, Moritz M. Daum and Stephanie Wermelinger	172

<i>Which Skills Debate Reaches the Public? Comparing Scientific Literature and Media Coverage of AI and LLM Skill Impacts (2022–2025)</i>	
Maud Reveilhac, Gerold Schneider, Bohdan Trembovelskyi, Nisha Yadav, Vlada Druta, Aurelia Tamò-Larrieux, Simon Mayer, Clement Guitton and Joshua C. Yang	180
<i>The Same Email, Signed Differently: Testing Negotiation Bias and Recommendation Stability in LLMs</i>	
Jasmin Heierli and Alexandre de Spindler	192
<i>A Bounded Coordination-Support Capability for Multi-Party Settings: Task-State Monitoring in Firefighter Incident Command</i>	
David Grünert, Barbara Morgenstern, Florian Peczinka, Dario Holenstein, Roland Brunner and Alexandre de Spindler	197
<i>Can Large Language Models Replace Statistical Software?</i>	
Yves Staudt	206
<i>Automated German Alt Text Generation for News Charts</i>	
Alessia Vannini, Julia Locher, Marianne Santaholma, Claudia Amsler, Jonas Oesch and Arzu Coltekin	215
2 Corpus Track	225
<i>A Dataset of Latin Etymologies Extracted from Wiktionary</i>	
Javier de Torres Méndez, Marco Carlo Passarotti, Giovanni Moretti, Francesco Mambrini and Matteo Pellegrini	226
<i>How Good is AI on Swiss Voting Booklets? A Multilingual OCR and Alignment Benchmark</i>	
Elina Stüssi and Jannis Vamvas	234