


MILLing 2024

Proceedings of the 2024 CLASP Conference on Multimodality and Interaction in Language Learning

Editors: Amy Qiu, Bill Noble, David Pagmar, Vladislav Maraev, and Nikolai Ilinykh



Gothenburg, Sweden and online
14–15 October 2024



CLASP centre for
linguistic theory
and studies in probability

CLASP Papers in Computational Linguistics, Volume 6
University of Gothenburg

CLASP Conference Proceedings, Volume 3
©2024 The Association for Computational Linguistics

Front-cover art: David Pagmar.

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

ISSN 2002-9764
ISBN 979-8-89176-163-6

Message from the organisers

We are happy to welcome you to the CLASP Conference on Multimodality and Interaction in Language Learning (MILLing 2024)! This volume consists of the archival papers presented at the MILLing conference held at the Department of Philosophy, Linguistics and Theory of Science (FLoV), University of Gothenburg on October 14–15, 2024. The purpose of this conference was to bring together researchers in linguistics and computational linguistics to discuss learning through linguistic interaction, from the perspectives of both human language acquisition and machine learning. The conference covers areas such as theoretical linguistics, experimental linguistics, pragmatics, computational linguistics, artificial intelligence, and cognitive science.

Recent transformer-based chat systems have impressed users with their ability to sustain coherent conversations on a wide range of topics. These models are created by fine-tuning large language models (LLMs) trained on massive text corpora. In spite of their impressive performance, these models are still fundamentally different from human speakers both in their linguistic ability and in the process by which they *learn* language. This leads to a number of questions about how the LLM learning procedure compares to language acquisition for human children: Can we design more data-efficient AI models? Audio, visual, and haptic information is available to humans both prior to and throughout language acquisition. Can such multi-modal information benefit AI language learning as well? Human language acquisition follows distinct stages, and different environmental inputs are available at different stages. Can curriculum learning benefit AI models? Pre-training + RLHF is fundamentally a batch learning procedure. Can AI models be designed to incorporate learning signals on the fly, as in online reinforcement learning? Explicit instruction is, colloquially, an important way that humans learn. Chat bots can follow prompt instructions, but prompts are ephemeral and the prompting window is finite. Can AI models be designed to learn from explicit instruction?

MILLing invited papers on topics from these and closely related areas, including (but not limited to): language acquisition (especially formal, statistical, experimental, and machine learning-based work); language learning through dialogue (in both humans and machines); multi-modality and figurativeness in language learning and dialogue; linguistic variation, adaptation, and audience design; low-resource and ecologically plausible language modelling (e.g., BabyLM); cognitive architectures for language learning; information state update in humans and machines; cognitive approaches to second language acquisition; dialogue systems for language learning; online, reinforcement and curriculum learning in NLP; atypical development and language learning; and ethical considerations in AI-assisted language learning.

Accepted papers and invited talks included topics ranging from socially intelligent agents, the relationship between probabilistic spaces and language learning, insights from neurodivergent language learners and people with language disorders, pre-linguistic communication, prosody, lexical meaning acquisition and adaptation, joint attention, language-as-action, spatial perspective coordination and more. The conference, and by extension these proceedings, is a discussion about these related topics and that examines various approaches and how they can mutually inform each other. The event included presentations of 10 accepted peer-reviewed papers, including 7 archival short papers and 2 archival long papers, 4 invited keynote talks, a panel discussion, and a poster session with 10 posters. We would like to thank all our contributors, programme committee members and volunteers, with special thanks to CLASP for organising the hybrid conference and the Swedish Research Council for funding CLASP.

Amy Qiu, Bill Noble, David Pagmar, Vladislav Maraev, and Nikolai Ilinykh

Gothenburg, Sweden

October 2024

Organising Committee

Program Chairs

Amy Qiu, University of Gothenburg, Sweden
Bill Noble, University of Gothenburg, Sweden
David Pagmar, University of Gothenburg, Sweden
Vladislav Maraev, University of Gothenburg, Sweden

Proceedings Chair

Nikolai Ilinykh, University of Gothenburg, Sweden

Local Arrangements Organisers

Susanna Myyry, University of Gothenburg, Sweden
David Pagmar, University of Gothenburg, Sweden
Amy Qiu, University of Gothenburg, Sweden

Programme Committee

Maxime Amblard	University of Lorraine
Fahima Ayub Khan	University of Gothenburg
Alexander Berman	University of Gothenburg
Ellen Breitholtz	University of Gothenburg
Stergios Chatzikyriakidis	University of Crete
Robin Cooper	University of Gothenburg
Simon Dobnik	University of Gothenburg
Eleni Gregoromichelaki	University of Gothenburg
Xudong Hong	Saarland University
Julian Hough	Swansea University
Christine Howes	University of Gothenburg
Nikolai Ilinykh	University of Gothenburg
Elisabetta Jezek	University of Pavia
Richard Johansson	Chalmers Technical University & University of Gothenburg
Ruth Kempson	King's College London
Nikhil Krishnaswamy	Colorado State University
Shalom Lappin	University of Gothenburg & King's College London
Staffan Larsson	University of Gothenburg
Sharid Loáiciga	University of Gothenburg
Andy Lücking	Goethe University Frankfurt
Chiara Mazzocconi	Aix Marseille University
Louise McNally	Universitat Pompeu Fabra
Gregory Mills	Kingston University
Massimo Poesio	Queen Mary University of London
Matthew Purver	Queen Mary University of London
James Pustejovsky	Brandeis University
Mehrnoosh Sadrzadeh	University College London
Asad Sayeed	University of Gothenburg
David Schlangen	University of Potsdam
Sabine Schulte im Walde	University of Stuttgart
Vidya Somashekarappa	University of Gothenburg
Carl Vogel	Trinity College Dublin
Sina Zarrieß	University of Bielefeld

Invited Speakers

Napoleon Katsos, University of Cambridge

Catherine Pelachaud, ISIR, UPMC

Charles Yang, University of Pennsylvania

Robin Cooper, University of Gothenburg

Invited talk: Napoleon Katsos

“Deficit”, “difficulty”, “difference”: perspectives into autistic people’s pragmatic skills and their implications for research methodology

It is widely reported that autistic people face pervasive challenges with producing and understanding pragmatics, i.e. context-dependent aspects of language. These are often attributed to challenges with mentalising, i.e. the ability to attribute the correct beliefs and intentions to other people. In this talk I will select influential papers from the past three decades of research in autism and language, each of which reveal a radically different perspective on the architecture of the linguistic system and on what it means to face challenges with linguistic competence (in our case, pragmatics). I will conclude that the recent perspective of neurodiversity implies a radical re-think of how we define pragmatics and how we assess the acquisition and processing of pragmatic competence.

Invited talk: Catherine Pelachaud

Multi-forms Adaptation for Socially Interactive Agents

Interacting with others enhances learning. Getting feedback on results, being encouraged and motivated . . . all help the learning process. During interaction, participants adapt to each other to show affiliation, group belongings, or to support social bonding. Adaptation can take place at different levels, through verbal alignment, imitation, and conversational strategies. Social resonance can also serve as a marker of adaptation. Socially Interactive Agents SIAs are virtual agents with a human-like appearance, capable of communicating verbally and nonverbally with their human interlocutors. In this talk, I will present our latest works aimed at endowing an SIA with various adaptive capabilities when interacting with its partners. The adaptation mechanisms are learned from human-human interaction data and evaluated by experimental studies involving human-agent interaction.

Invited talk: Charles Yang

Why language learning is not probabilistic

It seems harmless, and certainly mathematically convenient, to treat language learning as acquiring a probabilistic distribution over a space of linguistic patterns. The goal is to find or approximate an optimal hypothesis with respect to the data. Such is the mainstream machine learning approach, and the so-called Evaluation Procedure in generative grammar can be viewed as a particular instantiation. Despite having pursued it vigorously in my earlier work, I now believe this approach is wrong (and wrong-headed). On the one hand, language is not a zero sum game: even overwhelming presence of one linguistic form does not necessarily inhibit or penalise alternative forms. On the other, the grammar can be a partial function: there are inputs for which no output form is acceptable even though some will always be most highly valued in a probabilistic framework. The alternative is a theory of learning that does not even try to optimise but only satisfice. The coverage of the data only needs to be good enough up to a point; failure to do so may just result in the memorisation of the input – nothing in the cognitive system mandates generalisation under all circumstances. I will review the psychological and computational studies of the Tolerance Principle, a parameter-free learning theory that also appears operative beyond the domain of language.

Invited talk: Robin Cooper

Types in a Theory of Interactive Learning

In this talk I will present some CLASP research on using types in a theory of interactive learning. In the first part of the talk I will introduce the notion of type we have been using and how it relates to a general theory of action, including linguistic acts. In the second part of the talk I will present some work I have been doing with Staffan Larsson and Jonathan Ginzburg on how such a theory relates to communicative acts by prelinguistic children and how such communicative acts serve as a basis for the development of linguistic acts by children. In the third part of the talk, I will present some preliminary work with Staffan, Jonathan and Andy Lücking on how the kind of types we are using might relate to the approach to neural modelling that Chris Eliasmith and colleagues have been developing. While this work is still in the very early stages, the hope is that ultimately we can propose an explanatory account of interactive learning which is grounded in biologically plausible neural activity.

Table of Contents

<i>Critical Size Hypothesis: How Model Hyperparameters Correlate with Its Linguistic Abilities</i> Ekaterina Voloshina and Oleg Serikov	1
<i>INIKOL - Collocational Database for Learning Croatian as a Foreign Language</i> Goranka Blagus Bartolec, Gorana Duplančić Rogošić and Antonia Ordulj	8
<i>How Does an Adjective Sound Like? Exploring Audio Phrase Composition with Textual Embeddings</i> Saba Nazir and Mehrnoosh Sadrzadeh	13
<i>Learning through gesture: embodied repetitions in tandem interactions</i> Loulou Kosmala	19
<i>Towards Automated Game-Based Early Screening for Language Disorder</i> Hamdan Hamid Al-Ali, Elsa Soares, Goncalo Leal, Rita Valente, Nicole Agrela, Alexandra Marquis and Hanan Aldarmaki	26
<i>L2 Interaction in Heterogeneous Learner Groups during Content and Language Integrated Learning: The Experience of (removed for peer-review) and beyond</i> Julia Edeleva, Martin Neef, Jiaming Liu and Martin Scheidt	32
<i>Fifty shapes of BLiMP: syntactic learning curves in language models are not uniform, but sometimes unruly</i> Bastian Bunzeck and Sina Zarriß	39
<i>Not Just Semantics: Word Meaning Negotiation in Social Media and Spoken Interaction</i> Staffan Larsson, Jenny Myrendal and Bill Noble	56
<i>Toward Real Time Word Based Prosody Recognition</i> Alex Tilson and Frank Foerster	62

