


# LARP 2025

## Proceedings of the 2025 CLASP Conference on Language models and RePresentations

Editors: Nikolai Ilinykh, Erik Lagerstedt, Mattias Appelgren



Gothenburg, Sweden and online  
8–9 September 2025



**CLASP** centre for  
linguistic theory  
and studies in probability

CLASP Papers in Computational Linguistics, Volume 7  
University of Gothenburg

CLASP Conference Proceedings, Volume 4  
©2025 The Association for Computational Linguistics

Front-cover art: Erik Lagerstedt with the help of GPT-4.5.

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](mailto:acl@aclweb.org)

ISSN 2002-9764  
ISBN 979-8-89176-249-7

## Preface

We are delighted to welcome you to the CLASP conference on Language Models and RePresentations or **LARP 2025**! This volume consists of the archival papers presented at LARP, held at the Department of Philosophy, Linguistics and Theory of Science (FLoV), University of Gothenburg on September 8 – 9, 2025. The purpose of this conference was to bring together researchers in computational linguistics, artificial intelligence, and their intersections to discuss ideas on how language can be represented, and how computational language systems can both integrate neural (sub-symbolic) and symbolic representations. The conference covers areas such as computational linguistics, machine learning, artificial intelligence, natural language processing, and more.

The recent advances in language technology have been driven by the large language models (LLMs) built using transformers, large architectures, with representations built out of high dimensional feature spaces. These systems have been highly successful, however, much of human reasoning occurs on a symbolic level following the rules of logic, mathematics, or other systems. Classical AI was focused on symbolic systems, creating expert systems, planners, and search algorithms which manipulated systems, not tensors. There is a growing interest in the idea that these two methods can be combined in order to take advantage of the strengths of each. Many questions arise around these topics. How can neuro-symbolic architectures be created, what are the benefits and problems with them? Can these systems be used to create more explainable machine learning models? Can logical constraints imposed on neural networks increase both explainability, safety, and control over those models? Can automated reasoning provide human interpretable rationals for decisions? LARP invited papers on these topics and more. Accepted papers and invited talks included topics ranging from evaluating the reasoning capabilities of LLMs, bridging the gap between symbolic and neural approaches, abstractions for AI problem solving, to specific implementations of neuro-symbolic and reasoning systems. The conference, and by extension these proceedings, is a discussion about these related topics which examine various approaches and how they can mutually inform each other.

The event included 7 oral talks with presentations of 5 accepted peer-reviewed papers, including 1 archival short paper and 4 archival long papers. The event also had 3 invited keynote talks and a panel discussion. We would like to thank all our contributors, programme committee members, reviewers and volunteers, with special thanks to CLASP for organising the hybrid conference and the Swedish Research Council for funding CLASP.

Nikolai Ilinykh, Erik Lagerstedt, and Mattias Appelgren

Gothenburg, Sweden

September 2025

# Organizing Committee

## Program Chairs

Nikolai Ilinykh, University of Gothenburg, Sweden

Mattias Appelgren, University of Gothenburg, Sweden

Erik Lagerstedt, University of Gothenburg, Sweden

## Local Arrangements Chairs

Susanna Myyry, University of Gothenburg, Sweden

Mattias Appelgren, University of Gothenburg, Sweden

## Program Committee

Robin Cooper	University of Gothenburg
Alessandra Zarcone	Technische Hochschule Augsburg
Asad B. Sayeed	University of Gothenburg
Richard Johansson	Chalmers University of Technology and University of Gothenburg
Sharid Loáiciga	University of Gothenburg, Sweden
Maxime Amblard	Université de Lorraine
Bill Noble	Göteborg University and University of Gothenburg
Andy Lücking	Johann Wolfgang Goethe Universität Frankfurt am Main and Université Paris Diderot
Sandro Pezzelle	University of Amsterdam
Katrin Erk	University of Texas, Austin
Elisabetta Jezek	University of Pavia
Shalom Lappin	University of Gothenburg and Queen Mary University of London and King's College London
Nikhil Krishnaswamy	Colorado State University
Joakim Nivre	Uppsala University

**Keynote Talk**  
**PACE: Procedural Abstractions for Communicating  
Efficiently**

**Moa Johansson**

Chalmers University of Technology  
**2025-09-08 09:30:00 – Room: J222**

**Abstract:** A central but unresolved aspect of problem-solving in AI is the capability to introduce and use abstractions, something humans excel at. Work in cognitive science has demonstrated that humans tend towards higher levels of abstraction when engaged in collaborative task-oriented communication, enabling gradually shorter and more information-efficient utterances. In this talk, I will describe a neuro-symbolic method for introducing such abstractions called PACE. On the symbolic side, we draw on work from library learning in program synthesis for proposing abstractions. We combine this with neural methods for communication and reinforcement learning, via a novel use of bandit algorithms for controlling the exploration and exploitation trade-off in introducing new abstractions. Accepted for CogSci 2025 (oral), preprint: <https://arxiv.org/abs/2409.20120>

**Bio:** Moa Johansson is an Associate Professor in the Data Science and AI division at Chalmers University of Technology. She is interested in neuro-symbolic AI: the combination of neural machine learning methods and symbolic methods from e.g. theorem proving and program synthesis. Her group works on applications in maths and reasoning, cognitive science, and language.

# Keynote Talk

## On Retrieving & Reasoning LLMs: Myths, Merits, and How to Move Forward

**Dan Roth**

University of Pennsylvania and Oracle

2025-09-08 15:00:00 – Room: J222

**Abstract:** The rapid progress made over the last few years in generating linguistically coherent natural language has blurred, in the mind of many, the difference between natural language generation, understanding, knowledge retrieval and use, and the ability to reason with respect to the world. Nevertheless, reliably and consistently supporting high-level decisions that depend on natural language understanding and heterogenous information retrieval is still difficult, mostly, but not only, since most of these tasks are computationally more complex than language models can support. I will discuss some of the challenges underlying reasoning and information access and argue that we should exploit what LLMs do well while delegating responsibility to special purpose models and solvers for decision making. I will present some of our work in this space, focusing on supporting reasoning and information access via neuro-symbolic methods.

**Bio:** Dan Roth is the Eduardo D. Glandt Distinguished Professor at the Department of Computer and Information Science, University of Pennsylvania and the Chief AI Scientist at Oracle. Until June 2024 Dan was a VP/Distinguished Scientist at AWS AI. In his role at AWS Roth led over the last three years the scientific effort behind the first-generation Generative AI products from AWS, including Titan Models, Amazon Q efforts, and Bedrock, from inception until they became generally available. Dan is a Fellow of the AAAS, ACM, AAAI, and ACL. In 2017, Dan was awarded the John McCarthy Award; he was recognized for “for major conceptual and theoretical advances in the modeling of natural language understanding, machine learning, and reasoning”. He has published broadly in natural language processing, machine learning, knowledge representation and reasoning, and learning theory, was the Editor-in-Chief of the Journal of Artificial Intelligence Research (JAIR) and has served as a Program Chair and Conference Chair for the major conferences in his research areas. Roth has been involved in several startups; most recently he was a co-founder and chief scientist of NexLP, a startup that leverages the latest advances in Natural Language Processing, Cognitive Analytics, and Machine Learning in the legal and compliance domains. NexLP was acquired by Reveal. Dan received his B.A Summa cum laude in Mathematics from the Technion, Israel and his Ph.D. in Computer Science from Harvard University in 1995.

# Keynote Talk

## Reasoning with Large & Small Models: Bridging Symbolic and Neural Approaches

Vaishak Belle

University of Edinburgh

2025-09-09 10:00:00 – Room: J222

**Abstract:** This talk explores the intersection of large language models (LLMs) and reasoning systems, with a focus on addressing fundamental challenges in developing correct and reliable systems. We'll examine our work on augmenting LLMs with external "symbolic executors", creating hybrid architectures that leverage the strengths of both paradigms. The presentation will then talk about how LLMs represent and manipulate beliefs - standing for interactions with human or artificial users. We'll also discuss a few considerations for agentic pipelines, and how these sit with the broader paradigm of agent modelling, which has a long history in AI. We'll preface this development by first briefly reviewing the paradigm of neuro-symbolic AI, and emergent ideas such as loss functions and neural program induction.

**Bio:** Dr Vaishak Belle (he/him) is a Chancellor's Fellow and Reader at the School of Informatics, University of Edinburgh. He is an Alan Turing Institute Faculty Fellow, a Royal Society University Research Fellow, and a member of the RSE (Royal Society of Edinburgh) Young Academy of Scotland. He was previously at KU Leuven (Belgium), University of Toronto (Canada), Aachen University of Technology (Germany) and University of Trento (Italy). At the University of Edinburgh, he directs a research lab on artificial intelligence, specialising in the unification of logic and machine learning, with a recent emphasis on explainability and ethics. He has given research seminars at academic institutions such as MIT and Oxford, tutorials at AI conferences, and talks at venues such as Ars Electronica and the Samsung AI Forum. He has co-authored close to 120 peer-reviewed articles on AI, at venues such as IJCAI, UAI, AAAI, MLJ, AIJ, JAIR, AAMAS, and along with his co-authors, he has won the Microsoft best paper award at UAI, the Machine learning journal best student paper award at ECML-PKDD, and the Machine learning journal best student paper award at ILP. In 2014, he received a silver medal by the Kurt Goedel Society. He has served on the senior program committee/area chair of major AI conferences, co-chaired the ML track at KR, among others, and as PI and CoI secured a grant income of close to 8 million pounds. Recently, he has consulted with major banks on explainable AI and its impact in financial institutions.



## Table of Contents

<i>Simple Morphology, Complex Models: A Benchmark Study and Error Analysis of POS Tagging for Martinican Creole</i>	
Ludovic Mompelat .....	1
<i>EventHopNLI: A Functional Dataset for Systematically Diagnosing Logical Failures in LLM Temporal Reasoning</i>	
Ved Mathai and Janet B. Pierrehumbert .....	11
<i>Combining Information State Update, Harel Statecharts and LLMs for controllable and flexible Conversational AI</i>	
Vladislav Maraev, Alexander Berman and Staffan Larsson .....	28
<i>Towards Neuro-Symbolic Approaches for Referring Expression Generation</i>	
Manar Ali, Marika Sarzotti, Simeon Junker, Hendrik Buschmeier and Sina Zarrieß .....	38
<i>Extracting a Prototypical Argumentative Pattern in Financial Q&amp;As</i>	
Giulia D’Agostino, Michiel Van Der Meer and Chris Reed .....	51