

# Integrating Reasoning Systems for Trustworthy AI

Proceedings of the 4th Workshop on  
Logic and Practice of Programming (LPOP)

**(Preliminary Version)**

Held in conjunction with  
the 40th International Conference on Logic Programming (ICLP)  
Dallas, Texas, USA

October 13, 2024

Editors: Anil Nerode and Yanhong A. Liu

# Contents

<b>Preface</b>	<b>2</b>
<b>Program</b>	<b>5</b>
<b>Organization</b>	<b>7</b>
<b>Abstracts for Invited Talks</b>	<b>8</b>
How to Make Logics Neurosymbolic . . . . .	9
<i>Luc De Raedt</i>	
Psychoanalysis (and Therapy) of ChatGPT . . . . .	10
<i>Georg Gottlob</i>	
Title: Neuro-symbolic AI, or Are We Already There? . . . . .	11
<i>Henry Kautz</i>	
<b>Papers for Invited Talks</b>	<b>12</b>
Tools Are All You Need . . . . .	13
<i>Henry Kautz</i>	
<b>Refereed Position Papers</b>	<b>29</b>
Logic in the Age of AI . . . . .	30
<i>Yuri Gurevich</i>	
Rigorous Language Models for Trustworthy AI . . . . .	33
<i>Annie Liu</i>	
Generative AI as a Contributor to Joint Interactive Modeling . . . . .	37
<i>Joost Vennekens</i>	
Harnessing ASP and Its Extensions: Recent Applications and Role in Trustworthy AI . . . . .	40
<i>Giuseppe Mazzotta and Francesco Ricca</i>	
Beyond Deductive Datalog: How Datalog Is Used to Perform AI Tasks . . . . .	44
<i>Bernhard Scholz, Pavle Subotic and David Zhao</i>	
A Case Study on TSP: What to Optimize and How? . . . . .	47
<i>Martin Gebser</i>	
Declarative Ethics for AI Systems . . . . .	50
<i>John Hooker</i>	
Estimating Causal Quantities via Linear and Nonlinear Programming: Current Status, Challenges, and Future Directions . . . . .	54
<i>Ang Li</i>	

# Preface

Logical reasoning systems are essential for rigorous automatic reasoning. The focus of the 2024 Logic and Practice of Programming workshop is integrating reasoning systems for trustworthy AI, especially including integrating diverse models of programming with rules and constraints.

Trustworthy AI requires programming with rules and constraints for expressing and solving knowledge-intensive inference and combinatorial problems. A wide range of programming models have been proposed, including but not limited to the following, and essentially all of them require or support imperative programming for use in practical applications.

1. Classical first-order logic (FOL), not supporting transitive relations, with satisfiability (SAT) and satisfiability modulo theory (SMT) solvers
2. Deductive database (primarily Datalog) systems with fact-driven inference
3. Logic programming (dominantly Prolog) systems with goal-directed search, extended with sophisticated tabling and well-founded semantics
4. Answer set programming (ASP) systems, with sophisticated grounding and solving and stable model semantics
5. First-order logic (FOL) extended with inductive definitions (ID)
6. Constraint logic programming (CLP) extending Prolog systems with constraints
7. Constraint programming (CP), not supporting transitive relations, with backtracking, constraint propagation, local search, and more for solving
8. Mathematical programming (MP), not supporting transitive relations, with linear programming, nonlinear programming, and much more for solving
9. Co-inductive logic programming (s(ASP), s(CASP)) extending Prolog systems with goal-directed search for ASP and co-ASP solutions
10. Advanced knowledge representation (KR) with higher-order, objects, updates, defeasible reasoning, paraconsistency, uncertainty, and probability

Given any application problem—whether for planning or scheduling or regulatory compliance, requiring logical or probabilistic reasoning, or constraint satisfaction or optimization—how to best express and solve it using one or more of the models?

In recent years, AI systems built with large neural networks trained on massive data sets (such as GPT3 with 96 layers and 175-billion parameters on 570 GB of filtered data <https://arxiv.org/pdf/2005.14165.pdf>) have become increasingly capable in producing impressive outputs and beating humans in many applications. However, these systems may produce outputs that are not reliable, explainable, or aligned with intended uses.

The goal of the workshop is to bring together best people and best languages, tools, and ideas to discuss how to address these challenges with rigorous knowledge representation and reasoning, powerful and easy-to-use rule and constraint languages, and robust justifications and alignment checks. A wide variety of application problems will be used in the discussions. See here for some example application problems.

The workshop program consists of invited talks, presentations of position papers, and invited panels discussing key issues and future directions.

Potential workshop participants were invited to submit a position paper (1 or 2 pages in PDF format). Because we intend to bring together people from a diverse range of language and programming communities, it is essential that all talks be accessible to non-specialists.

The program committee invited attendees based on their position paper submissions and attempted to accommodate presentations in ways that fit with the broader organizational goal. Each submitted position paper, except for invited talks, was reviewed by at least three program committee members, and almost all accepted papers received at least two clear Accept's or even Strong Accept's, and none received any negative rating.

LPOP 2024 is a followup to three previous successful LPOP workshops held as part of the Federated Logic Conference (FLoC) in Oxford, UK in 2018; the ACM SIGPLAN conference on Systems, Programming, Languages, and Applications: Software for Humanity (SPLASH) in Chicago, USA, Virtual, in 2020; and by itself, Online, using Zoom, in 2022. LPOP 2018 focused on the integration of logic programming with imperative programming. LPOP 2020 broadened to the practical use of logic as a crosscutting discipline through many areas of computer science. LPOP 2022 focused on core high-level abstractions around sets and logic rules, to help bring them to the general practice of programming. LPOP 2024 focuses on integrating reasoning systems and their wide range of programming models for trustworthy AI.

LPOP 2024 includes invited talks by three distinguished researchers:

- Luc De Raedt (KU Leuven, Belgium) describes a way to develop Neurosymbolic AI by starting from a logic, adding a probabilistic interpretation, and then turning neural networks into 'neural predicates'.
- Georg Gottlob (Oxford University, UK) discusses shortcomings of large language models (LLMs), reasons for potential failures in such LLMs, and work to leverage such LLMs and mitigating the pitfalls.
- Henry Kautz (University of Virginia, US) explores giving LLMs the ability to use formal reasoning tools such as theorem provers and planners to achieve a base-level intelligence for further study of intelligence in general.

The program includes eight presentations by authors of contributed position papers, whose authors and titles are:

- Yuri Gurevich – Logic in the Age of AI
- Annie Liu – Rigorous Language Models for Trustworthy AI
- Joost Vennekens – Generative AI as a Contributor to Joint Interactive Modeling
- Giuseppe Mazzotta and Francesco Ricca – Harnessing ASP and Its Extensions: Recent Applications and Role in Trustworthy AI
- Bernhard Scholz, Pavle Subotic and David Zhao – Beyond Deductive Datalog: How Datalog Is Used to Perform AI Tasks
- Martin Gebser – A Case Study on TSP: What to Optimize and How?
- John Hooker – Declarative Ethics for AI Systems
- Ang Li – Estimating Causal Quantities via Linear and Nonlinear Programming: Current Status, Challenges, and Future Directions

Three invited panels are organized, with focused discussions on Integrating Logical Reasoning and LLMs for Trustworthy AI, chaired by Gopal Gupta (UT Dallas); Logic Systems and Optimizations for AI, chaired by David Warren (Stony Brook University and XSB Inc.); and Ethics, Uncertainty, and Roadmaps for AI, chaired by Anil Nerode (Cornell University).

The idea is to bring together experts from different communities to discuss views on how trustworthy AI systems might be developed in the coming years to take better advantage of different reasoning systems.

The overall organization, combining invited talks, paper presentations, and panels, is structured to encourage a deeper understanding of the various approaches and how they might mutually interact and benefit each other. We hope the participants enjoy the variety of talks and discussions!

We thank all LPOP program committee members for providing timely helpful and insightful reviews. Special thanks to David S. Warren for chairing the invited panel on Logic Systems and Optimizations for AI.

October 2024

Anil Nerode  
Y. Annie Liu

# Program

**Sunday October 13, 2024**

Displayed time zone: Central Time (US & Canada)

8:30–9:00 Breakfast

**9:00 Opening and Introduction**

Anil Nerode (Cornell University) and Annie Liu (Stony Brook University)

**9:15–10:30 Integrating Logic with LLMs for AI**

Chair: Gopal Gupta (UT Dallas)

**9:15 Invited Talk: Psychoanalysis (and Therapy) of ChatGPT**

Georg Gottlob (Oxford University)

**10:00 Logic in the Age of AI**

Yuri Gurevich (University of Michigan)

**10:15 Rigorous Language Models for Trustworthy AI**

Annie Liu (Stony Brook University)

10:30–11:00 Coffee Break

**11:00–12:30 Reasoning for AI and AI for Modeling**

Chair: Annie Liu (Stony Brook University)

**11:00 Invited Talk: Neuro-symbolic AI, or Are We Already There?**

Henry Kautz (University of Virginia)

**11:45 Generative AI as a Contributor to Joint Interactive Modeling**

Joost Vennekens (KU Leuven)

**12:00 Invited Panel: Integrating Logical Reasoning and LLMs for Trustworthy AI**

Georg Gottlob (Oxford University), Yuri Gurevich (University of Michigan), Annie Liu (Stony Brook University), Henry Kautz (University of Virginia), Joost Vennekens (KU Leuven), Chair: Gopal Gupta (UT Dallas)

12:30–13:30 Lunch

13:30–15:30 **Logic Systems and Optimizations for AI**  
 Chair: David Warren (Stony Brook University and XSB Inc.)

13:30 **Invited Talk: How to Make Logics Neurosymbolic**  
 Luc De Raedt (KU Leuven)

14:15 **Harnessing ASP and Its Extensions: Recent Applications and Role in Trustworthy AI**  
 Giuseppe Mazzotta and Francesco Ricca (University of Calabria)

14:30 **Beyond Deductive Datalog: How Datalog is used to perform AI tasks**  
 Bernhard Scholz (Sonic Research), Pavle Subotic (Sonic Research), and David Zhao (RelationalAI)

14:45 **A Case Study on TSP: What to Optimize and How?**  
 Martin Gebser (University of Klagenfurt)

15:00 **Invited Panel: Logic Systems and Optimizations for AI**  
 Luc De Raedt (KU Leuven), Giuseppe Mazzotta and Francesco Ricca (University of Calabria), Bernhard Scholz and Pavle Subotic (Sonic Research), David Zhao (RelationalAI), Martin Gebser (University of Klagenfurt), Chair: David Warren (Stony Brook University and XSB Inc.)

15:30–15:00 Coffee Break

16:00–17:30 **Ethics, Uncertainty, and Roadmaps for AI**  
 Chair: Anil Nerode (Cornell University)

16:00 **Declarative Ethics for AI Systems**  
 John Hooker (CMU)

16:15 **Estimating Causal Quantities via Linear and Nonlinear Programming: Current Status, Challenges, and Future Directions**  
 Ang Li (Florida State University)

16:30 **Combining Expressive Logic Programs with Machine Learning and Natural Language: Some Roadmap**  
 Invited Program Manager, Benjamin Grosof (DARPA)

17:00 **Invited Panel: Ethics, Uncertainty, and Roadmaps for AI**  
 John Hooker (CMU), Ang Li (Florida State University), Benjamin Grosof (DARPA), Chair: Anil Nerode (Cornell University)

17:30 **Closing**  
 Anil Nerode (Cornell University) and Annie Liu (Stony Brook University)

# Organization

## Chairs

Anil Nerode, Cornell University, US  
Annie Liu, Stony Brook University, US

## Program Committee

Martin Gebser, University of Klagenfurth, Austria  
Michael Gelfond, Texas Tech University, US  
Benjamin Grosof, DARPA, US  
Gopal Gupta, UT Dallas, US  
Michael Kifer, Stony Brook University, US  
Marta Kwiatkowska, University of Oxford, UK  
Fabrizio Riguzzi, University of Ferrara, Italy  
Joost Vennekens, KU Leuven, Belgium  
Toby Walsh, University of New South Wales, Australia  
Jan Wielemaker, CWI, The Netherlands  
Roland Yap, National University of Singapore, Singapore

## LPOP Website

<https://lpop.cs.stonybrook.edu/>