

Joint Workshop on Statistics and Knowledge Integration for Logic, Learning, Ethical Decisions, and LLMs (SKILLED-LLMs 2026)

Co-located with FLoC 2026

Lisbon, Portugal, July 18, 2026

Preface

Ha Thanh Nguyen <i>NII</i> Tokyo, Japan	Francesca Toni <i>Imperial College</i> London, UK	Kostas Stathis <i>RHUL</i> London, UK	Ken Satoh <i>ROIS-DS</i> Tokyo, Japan	Randy Goebel <i>University of Alberta</i> Alberta, Canada	Francesco Chiariello <i>Inria & ENS — PSL</i> France
---	---	---	---	---	--

Yves Lesperance <i>York University</i> Toronto, Canada	Matteo Magnini <i>Luxembourg University</i> Luxembourg	Federico Sabbatini <i>Bologna University</i> Bologna, Italy	Elena Umili <i>Sapienza University</i> Rome, Italy	Nourhan Ehab <i>GUC</i> Cairo, Egypt	Mervat AbuElkheir <i>GUC</i> Cairo, Egypt
--	--	---	--	--	---

We are pleased to present the proceedings of the Joint Workshop on Statistics and Knowledge Integration for Logic, Learning, Ethical Decisions, and LLMs (SKILLED-LLMs 2026), affiliated with the 23rd International Conference on Principles of Knowledge Representation and Reasoning (KR 2026), co-located with the 9th Federated Logic Conference (FLoC 2026), and held in Lisbon, Portugal, on July 18, 2026.

Reasoning is an essential component of human intelligence as it plays a fundamental role in our ability to think critically, support responsible decisions, and solve challenging problems. Traditionally, artificial intelligence (AI) has addressed reasoning in the context of logic-based representations of knowledge. However, the recent leap forward in natural language processing, with the emergence of language models based on transformers, is hinting at the possibility that these models exhibit reasoning abilities, particularly as they grow in size and are trained on more data.

The goal of the SKILLED-LLMs workshop was to create a platform for researchers from different disciplines and AI perspectives to explore approaches and techniques for reconciling reasoning between deep learning and logic-based representations. This includes integrating knowledge-representation-style reasoning with deep learning models to build neuro-symbolic systems, as well as analysing the reasoning abilities of large language models and formalising the kinds of reasoning they carry out. Furthermore, as language models become increasingly prevalent in society, it is crucial to address the ethical implications of their deployment and use, including bias, fairness, transparency, and accountability. Finally, the workshop was open to any approach that combines data-driven techniques with knowledge representation and reasoning, whether based on deep learning, statistics, or other paradigms.

The SKILLED-LLMs workshop merged three separate workshops proposals submitted to KR 2026:

- NeLaMKRR: Next-Generation Language Models for Knowledge Representation and Reasoning
- SKILL: Statistics and Knowledge Integration for Logic and Learning
- ReLLM: Reasoning and Ethics in Large Language Models

By merging these three communities, we aimed at fostering interdisciplinary collaboration and create a comprehensive forum for discussing reasoning, ethics, and knowledge integration in AI systems.

The central themes of the workshop revolved around the intersections of language models, knowledge representation, reasoning, ethics, and statistics. We encouraged submissions discussing novel techniques, approaches, and innovative ideas related to these topics. The list of topics of interest included, but was not limited to:

- Language models' reasoning abilities
- KR-style reasoning in LLMs
- Ethical deployment and bias mitigation

- Integration of symbolic and statistical methods
- Neurosymbolic approaches
- Domain-specific applications (medicine, law, science)
- Chain of thought and prompting studies
- Evaluation of reasoning accuracy

The workshop paper submission deadline was fixed to May 14, 2026. Acceptance and reject notifications were sent on May 28, 2026. The Organizing Committee received nineteen submissions. Submitted papers were peer-reviewed by at least two expert scholars of the field. Feedback were based on technical quality, relevance, originality, significance, and clarity. Sixteen papers were accepted for presentation as short and regular papers. These papers were published in the present volume. The maximum length for regular and short papers was fixed to 12 and 6 pages, respectively, excluding references, generative AI declarations, and bibliography.

We would like to sincerely thank our Program Committee members for their precious work:

- Damiano Azzolini, University of Ferrara, Italy
- Daniel Sonntag, German Research Center for Artificial Intelligence (DFKI), Germany
- Fabrizio Riguzzi, University of Ferrara, Italy
- Gabriel Freedman, Imperial College London, United Kingdom
- Gianluca Cima, Sapienza University of Rome, Italy
- Haythem Ismail, German University in Cairo, Egypt
- Justin Brody, Franklin & Marshall College, USA
- Kislay Raj, Dublin City University, Ireland
- Laura Papi, Sapienza University of Rome, Italy
- Leon van der Torre, University of Luxembourg, Luxembourg
- Luke Dickens, University College London, UK
- Marco Console, Sapienza University of Rome, Italy
- María Navas-Loro, Universidad Politécnica de Madrid, Spain
- Matteo Mancanelli, Sapienza University of Rome, Italy
- May Myo Zin, Center for Juris-Informatics (ROIS-DS), Japan
- Mayar Abdelwahab, German University in Cairo, Egypt
- Neil Yorke-Smith, TU Delft, Netherlands
- Pedro Zuidberg Dos Martires, Örebro University, Sweden
- Pietro Totis, J.P. Morgan, Madrid, Spain
- Roberta Calegari, University of Bologna, Italy
- Sabine Wehnert, Ruhr University Bochum, Germany
- Siraj Munir, University of Urbino, Italy
- Thi-Hai-Yen Vuong, VNU University of Engineering and Technology, Vietnam
- Tilman Hinnerichs, TU Delft, Netherlands
- Vince Trencsenyi, Royal Holloway University of London, United Kingdom
- Vu Tran, Japan Advanced Institute of Science and Technology, Japan
- Wachara Funwacharakorn, Center for Juris-Informatics (ROIS-DS), Japan

July 2026

The Organising Committee

Ha Thanh Nguyen, NII, Tokyo, Japan
 Francesca Toni, Imperial College, London, UK
 Kostas Stathis, RHUL, London, UK
 Ken Satoh, ROIS-DS, Tokyo, Japan
 Randy Goebel, University of Alberta, Alberta, Canada
 Francesco Chiariello, Inria & ENS — PSL, France
 Yves Lesperance, York University, Toronto, Canada
 Matteo Magnini, Luxembourg University, Luxembourg
 Federico Sabbatini, Bologna University, Bologna, Italy
 Elena Umili, Sapienza University, Rome, Italy
 Nourhan Ehab, GUC, Cairo, Egypt
 Mervat AbuElkheir, GUC, Cairo, Egypt