Preface

Marco Sälzer[1], François Schwarzentruber[2] and Nicolas Troquard[3]

[1]RPTU University of Kaiserslautern-Landau, Germany
[2]ENS de Lyon, CNRS, Université Claude Bernard Lyon 1, Inria, LIP, UMR 5668, 69342, Lyon, France
[3]Gran Sasso Science Institute (GSSI), Viale F. Crispi, 7 – 67100 L'Aquila, Italy

Abstract

This volume contains papers presented at the *Workshop on Logical Methods for Neural Network Analysis* (LOGICNN'26), held in Lisbon, Portugal, on 25 July 2026. The workshop was affiliated with the *23rd International Conference on Principles of Knowledge Representation and Reasoning* (KR 2026), which was part of the *9th Federated Logic Conference* (FLoC 2026). The one-day programme featured three sessions of contributed talks covering different aspects of the interaction between logic and neural networks, as well as an invited talk and an invited tutorial. Further information is available at https://perso.ens-lyon.fr/francois.schwarzentruber/research/events/logicnn2026/.

Introduction

Neural networks are at the heart of modern machine learning and underlie many architectures used in contexts such as vision, language, and graph-structured data. They power safety-critical applications ranging from steering autonomous vehicles to supporting medical diagnostics, but their behaviour is notoriously difficult to predict, explain, or verify. While empirical evaluation is indispensable, it necessarily considers only finitely many inputs and cannot by itself establish that a model satisfies a property on all valid inputs. This inherent limitation has contributed to increasing interest in using formal methods, and logic in particular, as fundamental tools for analysing neural network models , specifying desired or unwanted behaviour, and deriving rigorous guarantees .

The role of logic in this context is not one-dimensional. Logical characterisations can be used to identify which properties a neural architecture can express, compare seemingly different models, and expose limitations imposed by their design. This viewpoint has been especially fruitful for graph neural networks , but has also been extensively investigated for other models, such as recurrent and transformer-based architectures . At another level, logical specifications turn questions about reachability, robustness, and safety into formal reasoning problems for which automated tools can search for proofs or counterexamples . The interaction also runs in the opposite direction: symbolic knowledge can be incorporated into learning , while, in suitable settings, logical rules can be extracted from trained networks , supporting formal reasoning and human-understandable explanations.

The papers contributed to LOGICNN'26 reflected this breadth. They addressed the logical characterisation of graph-based and recurrent models, sound neural network verification and verified neural computation, and translations of neural networks into logical formalisms. Together with the invited talk on recurrent GNNs and the tutorial on Marabou, the programme presented these as complementary approaches to the formal analysis of neural networks.

LOGICNN'26: Workshop on Logical Methods for Neural Network Analysis, July 25, 2026, Lisbon, Portugal
marco.saelzer@rptu.de (M. Sälzer); francois.schwarzentruber@ens-lyon.fr (F. Schwarzentruber); nicolas.troquard@gssi.it (N. Troquard)
https://marcosaelzer.github.io/ (M. Sälzer); https://perso.ens-lyon.fr/francois.schwarzentruber (F. Schwarzentruber); https://troquard.bitbucket.io/ (N. Troquard)
0000-0002-8012-5465 (M. Sälzer); 0000-0002-1228-4333 (F. Schwarzentruber); 0000-0002-5763-6080 (N. Troquard)

Objectives

The LOGICNN'26 workshop aimed to bring together researchers working at the intersection of logic and neural-network-based machine learning. It was intended as a forum for exchanging recent results, presenting work in progress, and identifying research questions for which logical methods can contribute to the understanding and analysis of neural networks. The call for papers was organised around four main themes:

- *Expressivity and translation.* Can logical formalisms characterise the behaviour of neural networks? This includes translations of neural architectures, such as CNNs, GNNs, RNNs, and transformers, into logical specifications, as well as translations in the opposite direction.
- *Safety verification.* Can logical methods provide formal guarantees that neural networks satisfy safety-critical properties?
- *Logical interpretation.* Can logic support the construction of human-understandable and faithful explanations of neural network behaviour?
- *Decidability and computational complexity.* Which problems arising in neural network analysis are decidable, and what is their computational complexity?

Submissions and Programme

The submission deadline was 27 April 2026, and authors were notified on 29 May 2026. The workshop received nine submissions: four regular papers, two short papers, three posters, including those of previously published work. Each submission received two reviews from members of the 22-person Programme Committee. All nine submissions were accepted for presentation and are included in the present volume.

The one-day programme comprised three thematic sessions on expressiveness, verification, and translations from neural networks into logic. Each contributed talk consisted of a 15-minute presentation followed by 5 minutes for discussion. The programme further included a 60-minute invited talk by Martin Grohe (RWTH Aachen University) and a 60-minute invited tutorial by Omri Isac (Hebrew University of Jerusalem).

Invited talk: Martin Grohe
Logic and the Power of Recurrent Graph Neural Networks

Session 1: Expressiveness
Expressive Power of Graph Transformers via Logic
Towards Understanding the Expressive Power of GNNs with Global Readout
An Algebraic Characterization of Local Weisfeiler–Leman
Words and Temporal Graphs: Comparing the Expressive Power of State Space Models and Recurrent Neural Networks

---

Invited tutorial: Omri Isac
Verification of DNNs with Marabou

Session 2: Verification
Towards Continuous Constraint Programming for Sound Neural Network Verification
IsaGrad: Verified Automatic Differentiation over Computational Graphs in Imperative HOL

Session 3: From Neural Networks to Logic
Neural Networks as Fuzzy Logic Formulas
Neural Networks into Łukasiewicz Logic, with Applications to Formal Verification
From MLPs to Logic: An End-to-End Neuro-Symbolic Compilation Pipeline for Explainable Safety-Critical Medical AI

Acknowledgements

We thank the authors and invited speakers for their contributions, the members of the Programme Committee for their careful and timely reviews, and all participants for the discussions. We are also grateful to the FLoC 2026 organisers for their support in preparing and hosting the workshop.

August 2026
Marco Sälzer, François Schwarzentruber, and Nicolas Troquard
LOGICNN'26 workshop co-chairs

Programme Committee