Mihaela Cătălina Stoian

Department of Computer Science – University of Oxford – Oxford, UK
mihaela.stoian@cs.ox.ac.uk • mihaela-stoian.github.io

Research Interests

Topics
Neuro-Symbolic AI
Knowledge-Aided Machine Learning
Deep Generative Modelling

Overview
My research focuses on developing neuro-symbolic AI systems that integrate background knowledge requirements into neural networks during training and ensure that these requirements are satisfied. A key aspect of my work involves creating such knowledge-aided machine learning frameworks that are both accessible and practical for practitioners working on real-world applications, e.g., synthesising tabular data and safe and trustworthy autonomous driving.

Education

University of Oxford **Oxford, UK**
DPhil in Computer Science 2021–present
Supervisor: Prof. Thomas Lukasiewicz
Topic: Refining Deep Neural Networks with Background Knowledge
Research areas: Neuro-symbolic AI, Generative Modelling
I developed the first framework that ensures deep generative models for tabular data synthesis are inherently compliant with background knowledge requirements, which are expressed using linear constraints, and, in my most recent work, using quantifier-free linear real arithmetic.

The University of Edinburgh **Edinburgh, UK**
Master of Informatics with Honours, First Class 2014–2019
Supervisor: Prof. Sharon Goldwater
Master's thesis on speech-to-text machine translation

Awards

EPSRC Scholarship for Doctoral Studies
University of Oxford 2021–2025

Women in Quant Finance Grant
G-Research 2025

Travel Grant
IJCAI-AIJ 2024

Travel for Research Grant
St Hilda's College, University of Oxford 2023, 2024

Publications

[1] **M. C. Stoian** and E. Giunchiglia. Beyond the Convexity Assumption: Realistic Tabular Data Generation under Quantifier-Free Real Linear Constraints. In Proc. of ICLR 2025.

[2] **M. C. Stoian**, S. Dyrmishi, M. Cordy, T. Lukasiewicz, and E. Giunchiglia. How Realistic Is Your Synthetic Data? Constraining Deep Generative Models for Tabular Data. In Proc. of ICLR 2024.

[3] **M. C. Stoian**. Deep Learning with Requirements in the Real World. In Proc. of IJCAI, Doctoral Consortium 2024.

[4] S. Dyrmishi, **M. C. Stoian**, E. Giunchiglia, M. Cordy. Deep generative models as an adversarial attack strategy for tabular machine learning. In Proc. of International Conference on Machine Learning and Cybernetics 2024.

[5] **M. C. Stoian**, A. Tatomir, T. Lukasiewicz, and E. Giunchiglia. PiShield: A PyTorch Package for Learning with Requirements. In Proc. of IJCAI 2024.

[6] E. Giunchiglia, A. Tatomir, **M. C. Stoian**, T. Lukasiewicz. CCN+: A neuro-symbolic framework for deep learning with requirements. In International Journal of Approximate Reasoning, Volume 171 (2024).

[7] S. Khan, I. Teeti, R. J. Alitappeh, **M. C. Stoian**, E. Giunchiglia, G. Singh, A. Bradley, F. Cuzzolin. ROAD-Waymo: Action Awareness at Scale for Autonomous Driving. arXiv preprint 2411.01683, 2024.

[8] **M. C. Stoian**, E. Giunchiglia, and T. Lukasiewicz. Exploiting T-norms for Deep Learning in Autonomous Driving. In Proc. of NeSy 2023.

[9] E. Giunchiglia, **M. C. Stoian**, S. Khan, F. Cuzzolin, and T. Lukasiewicz. ROAD-R: The Autonomous Driving Dataset with Logical Requirements. In Machine Learning, Volume 112 (2023). (best paper award at the AI4AD workshop hosted by IJCAI 2022 and best student paper prize at IJCLR 2022)

[10] E. Giunchiglia, **M. C. Stoian**, T. Lukasiewicz. Deep Learning with Logical Constraints. In Proc. of IJCAI, 2022.

[11] **M. C. Stoian**, T. Cavallari. Recurrently Estimating Reflective Symmetry Planes from Partial Pointclouds. In CVPR Workshop on 3D Vision and Robotics 2021.

[12] **M. C. Stoian**, S. Bansal, and S. Goldwater. Analyzing ASR pretraining for low-resource speech-to-text translation. In Proc. of ICASSP 2020.

Work Experience

FiveAI **Oxford, UK**
Research Intern 2020–2021
Supervisor: Dr. Tommaso Cavallari
Paper: M. C. Stoian, T. Cavallari. Recurrently Estimating Reflective Symmetry Planes from Partial Pointclouds. In CVPR Workshop on 3D Vision and Robotics, 2021.
Project recognised with a patent application.
I developed a novel method to estimate planar reflective symmetries that efficiently handles 3D inputs by slicing the data along the height dimension and passing it sequentially to a 2D convolutional recurrent regression scheme.

The University of Edinburgh **Edinburgh, UK**
Research Assistant 2019
Supervisor: Prof. Sharon Goldwater
Paper: M. C. Stoian, S. Bansal, and S. Goldwater. Analyzing ASR pretraining for low-resource speech-to-text translation. In Proc. of ICASSP 2020.

I showed that (i) combining automatic speech recognition (ASR) pretraining with data augmentation on the target language improves speech-to-text translation (AST) performance, (ii) the word error rate of the pretrained ASR models is a suitable direct predictor of AST performance.

ETH Zurich **Zurich, Switzerland**
Student Summer Research Fellow 2018
Supervisors: Prof. Martin Vechev, Assistant Prof. Dana Drachsler Cohen
Project: Program Behaviour Synthesis for Programming Protocol-Independent Packet Processors

The University of Edinburgh **Edinburgh, UK**
Research Intern 2017
Supervisor: Prof. Kousha Etessami
Project: Implementing algorithm for computing reachability in Branching Markov Processes

The University of Edinburgh **Edinburgh, UK**
Research Intern 2016
Supervisor: Prof. Sharon Goldwater
Project: Evaluating speech features from transcriptions using the ABX minimal pair criterion

Talks

Dagstuhl Seminar on Logic and Neural Networks	February 2025
IJCAI Doctoral Consortium	August 2024
Oxford Computer Science Conference	June 2024
University of Luxembourg	June 2024
Sony AI	April 2024

Teaching

The University of Edinburgh
Discrete Mathematics and Mathematical Reasoning 2018–2019

The University of Edinburgh
Algorithms, Data Structures and Learning 2017–2019

The University of Edinburgh
Processing Formal and Natural Languages 2017–2019

Other Experience

Supervisor
Co-supervising two current Master's students

Organiser
ROAD++: The Third Workshop & Challenge, hosted by ECCV 2024
ROAD-R: The Road Event Detection with Requirements Challenge, hosted by NeurIPS 2023
ROAD++: The Second Workshop & Challenge, hosted by ICCV 2023

Reviewer
Conferences: NeurIPS, ICLR, IJCAI, ICML, NeSy, ICPR
Workshops: ROAD++ (ICCV 2023), RepL4NLP (ACL 2022), NeSy-GeMs (ICLR 2023)
Journals: Machine Learning