Mihaela Cătălina Stoian

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

Education

University of Oxford **Oxford, UK**
DPhil in Computer Science 2021–present

PhD Thesis: Refining Deep Neural Networks with Background Knowledge for Real-World Applications
Supervisor: Prof. Thomas Lukasiewicz
Research areas: Deep learning, Generative modelling, Autonomous driving, Neuro-symbolic AI

I developed the first frameworks that ensure deep generative models for tabular data synthesis are inherently compliant with background knowledge requirements, which are expressed using linear constraints and disjunctions of linear inequalities that capture non-convex and even disconnected spaces.

The University of Edinburgh **Edinburgh, UK**
Master of Informatics with Honours, First Class 2014–2019

Master's Thesis: Exploring methods for improving extrinsic performance in zero-resource settings
Bachelor's Thesis: Exploring multilingual bottleneck features for zero-resource speech processing
Supervisor: Prof. Sharon Goldwater
Research areas: NLP, Speech-to-text machine translation, Speech processing

Awards

Oxford PhD Runner-up Prize awarded by G-Research 2025
EPSRC Scholarship for Doctoral Studies awarded by University of Oxford 2021–2025
IJCAI Grant awarded by IJCAI-AIJ 2024
Conference Travel Grant awarded by St Hilda's College, University of Oxford 2023, 2024
Best paper award at the AI4AD workshop at IJCAI 2022
Best student paper prize at IJCLR 2022

Work Experience

FiveAI **Oxford, UK**
Research intern in computer vision and autonomous vehicles 2020–2021
Paper: M. C. Stoian, T. Cavallari. Recurrently Estimating Reflective Symmetry Planes from Partial Pointclouds. In CVPR Workshop on 3D Vision and Robotics, 2021.
Supervisor: Dr. Tommaso Cavallari

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

ETH Zurich **Zurich, Switzerland**
Summer research fellow in network verification and program synthesis 2018
Project: Program Behaviour Synthesis for Programming Protocol-Independent Packet Processors
Supervisors: Prof. Martin Vechev, Assistant Prof. Dana Drachsler Cohen

Selected 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] S. Dyrmishi, **M. C. Stoian**, E. Giunchiglia, M. Cordy. Deep generative models as an adversarial attack strategy for tabular machine learning. In Proc. of ICMLC 2024.

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

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

[6] E. Giunchiglia, **M. C. Stoian**, S. Khan, F. Cuzzolin, and T. Lukasiewicz. ROAD-R: The Autonomous Driving Dataset with Logical Requirements. In Machine Learning, Vol. 112, 2023.

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

Selected Invited Presentations

Civil Service Leadership Group meeting, Imperial College London 2025
Presented my work (one of only four projects selected for the event) to senior UK civil servants, including the Head of the Civil Service.

Dagstuhl Seminar on Logic and Neural Networks 2025
Invitation-only event gathering experts in the field.

Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg 2024
Reading group of the Security, Reasoning and Validation research group.

Sony AI 2024
Barcelona and Tokyo reading groups.

Leadership

Co-supervisor
Current students (Vienna University of Technology): Alina Godun (Master's), Luka Pejic (Bachelor's)
Past students: Shinde Lee (Master's, Imperial College London)

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

Technical Skills

Programming Languages: Python (proficient), C++, Java, Shell (familiar)
Deep Learning: PyTorch, TensorFlow, Hugging Face Transformers, PyTorch Lightning
Data Science & Machine Learning: Scikit-learn, NumPy, SciPy, Matplotlib, Seaborn, Pandas
Experiment Tracking: Weights & Biases, TensorBoard
Tools & Infrastructure: Git, Docker, Slurm

Service

Reviewer for conferences: NeurIPS, ICLR, IJCAI, ICML, NeSy
Reviewer for the Machine Learning journal

Language Skills

English (proficient), Romanian (native), German (elementary)