

Autoformalization & Mathematical Reasoning

Lan Zhang, Marco Valentino, Jordan Meadows, André Freitas

<https://sites.google.com/view/autoformalization-tutorial>

Recent advances in large language models, formal reasoning, theorem proving, and agentic AI are rapidly changing what can be augmented and automated in mathematics. Yet mathematical reasoning is considerably richer than producing correct final answers: it involves understanding mathematical discourse, constructing and revising arguments, moving between informal and formal representations, exploring alternative proof strategies, verifying intermediate claims, and ultimately supporting mathematical discovery. This tutorial provides a structured introduction to the emerging field of automated mathematical reasoning, with a particular focus on autoformalization and neuro-symbolic and agentic approaches. We first examine mathematical reasoning from linguistic and conceptual perspectives, decomposing the spectrum from informal mathematical discourse to fully formal proofs. We then introduce the principal computational mechanisms underlying modern systems, including formal representations, autoformalization, theorem provers, critics and proxy evaluators, refinement loops, and complex agentic architectures. The tutorial connects these foundations with practical system-building, evaluation methodologies, current state-of-the-art results, and demonstrations of existing frameworks. We further examine how mathematical reasoning methods can extend toward physics, where mathematical arguments must interact with physical models, assumptions, approximations, and scientific knowledge. Finally, we use ambitious mathematical and scientific problems to identify the capabilities still required to move from benchmark solving toward systems that can participate meaningfully in verification, reasoning, and discovery.

Lan Zhang

email: lan.zhang-6@postgrad.manchester.ac.uk

website: <https://lanzhang128.github.io/>

Lan is a PhD student in the Department of Computer Science in University of Manchester. His research interest mainly focuses on autoformalization and formal mathematical reasoning. He has published papers at major NLP conferences, including ACL, EMNLP, and NAACL. His contributions to autoformalization were recognized with the Best Resource Paper Award at EMNLP 2025, highlighting the impact of his work on formal mathematical reasoning and verification.

Marco Valentino

email: m.valentino@sheffield.ac.uk

website: <https://www.marcovalentino.net/>

Marco is a Lecturer in Artificial Intelligence and Applications of AI in the Natural Language Processing (NLP) group at the University of Sheffield. Prior to Sheffield, he was a member of the Neuro-Symbolic AI Group at the Idiap Research Institute in Switzerland, and obtained a PhD in Computer Science from the University of Manchester. His research focuses on developing AI systems that can use explanation as a core mechanism for learning and reasoning, investigating the integration of neural and symbolic AI methods. Moreover, he is interested in developing methodologies to interpret, control, and evaluate Large Language Models (LLMs), with a focus on disentangling knowledge acquisition from abstract logical reasoning, and enabling out-of-distribution, out-of-domain generalisation. His research on neuro-symbolic NLP received a best resource paper award at EMNLP 2025 and an outstanding paper award at EMNLP 2024.

Jordan Meadows

email: jmeadows@idiap.ch

Jordan Meadows is a Postdoctoral Researcher at Idiap Research Institute working on natural language processing, neuro-symbolic AI, and AI for Science. He received his PhD in Computer Science from the University of Manchester, where he studied mathematical language processing and the reasoning capabilities of language models. His research focuses broadly on mathematical and scientific reasoning, with particular interest in combining language models with symbolic methods to make AI systems more reliable, interpretable, and verifiable. His work has been published across major NLP and AI venues, and he has presented research on mathematical reasoning and the autoformalisation of science to both academic and interdisciplinary audiences.

André Freitas

email: andre.freitas@idiap.ch

website: <https://www.andrefreitas.net/>

André Freitas leads the Neuro-symbolic AI Lab at the University of Manchester and IDIAP Research Institute. His main research interests are on enabling the development of AI methods to support abstract, flexible and controlled reasoning in order to support AI-augmented scientific discovery. In particular, he investigates how the combination of neural and symbolic data representation paradigms can deliver better models of inference. He is an active contributor to the main conferences and journals in the AI/Natural Language Processing (NLP) interface (AAAI, NeurIPS, ACL, EMNLP, EACL, COLING, TACL, Computational Linguistics), with over 100 peer-reviewed publications. He contributed to the organization of MathNLP at EMNLP 2022 and LREC-COLING 2024. André participated in 7 tutorials, and coorganized 1 conference and 6 workshops.