Nicolae Cudlenco

PhD Candidate in Computer Science
Institute of Mathematics „Simion Stoilow" of the Romanian Academy
nicolae.cudlenco@gmail.com scholar.google.com/citations?user=WB0Fx1gAAAAJ
linkedin.com/in/nicolae-cudlenco orcid.org/0000-0001-6547-3659
github.com/ncudlenco



RESEARCH PROFILE

PhD candidate in Computer Science at the School of Advanced Studies of the Romanian Academy, Institute of Mathematics „Simion Stoilow", supervised by Prof. Dr. Marius Leordeanu. My research studies what multimodal models understand about the physical world, and how **explicit, structured representations** can make that understanding measurable and explainable. It began with brain signals, showing that EEG signals can improve automatic visual recognition, and led to the **Graph of Events in Space and Time (GEST)**, a hierarchical event representation shared by vision and language. I built **GEST-Engine**, which executes such graphs, written by hand or by tool-constrained LLM agents, as multi-actor videos with exact per-frame ground truth, and I use it to probe and train video and vision-language models. In parallel, as Lead ML Architect at BÜCHI Labortechnik AG, I develop self-supervised pretraining for near-infrared spectroscopy, towards an NIR foundation model, and own the cloud and MLOps architecture of the company's NIR modelling platform.

EDUCATION

PhD in Computer Science Nov 2018 to present (defense 26 November 2026)
Institute of Mathematics „Simion Stoilow" of the Romanian Academy, SCOSAAR, Bucharest, Romania
- Thesis: „Understanding Visual Understanding: From Brain Signals to Explainable Multimodal Representations"
- Supervisor: Prof. Dr. Marius Leordeanu

MSc in Artificial Intelligence 2016 to 2018
University Politehnica of Bucharest, Faculty of Automatic Control and Computer Science. Final grade 10/10
- Thesis: „Decoding bioelectric signals of the human body"

BSc in Computer Science 2012 to 2016
University Politehnica of Bucharest, Faculty of Automatic Control and Computer Science. Final grade 10/10
- Thesis: „Optimized Feature Selection for eHealth Data"

PUBLICATIONS

1. **N. Cudlenco**, N. Popescu, and M. Leordeanu, „Reading into the mind's eye: Boosting automatic visual recognition with EEG signals", Neurocomputing, vol. 386, pp. 281-292, 2020. Web of Science, IF 4.438.
   URL: https://www.sciencedirect.com/science/article/pii/S0925231219317850
2. M. Masala, **N. Cudlenco**, T. Rebedea, and M. Leordeanu, „Explaining vision and language through graphs of events in space and time", in 2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW), pp. 2826-2831, IEEE, 2023.
   URL: https://openaccess.thecvf.com/ICCV2023_workshops/CLVL, https://arxiv.org/abs/2309.08612
3. **N. Cudlenco**, M. Masala, and M. Leordeanu, „[Tiny Paper] GEST-Engine: Controllable Multi-Actor Video Synthesis with Perfect Spatiotemporal Annotations", in ICLR 2026, 2nd Workshop on World Models: Understanding, Modelling and Scaling, 2026.
   URL: https://openreview.net/forum?id=uUofPYVMZH
4. **N. Cudlenco**, M. Masala, and M. Leordeanu, „The GEST-Engine: From Event Graphs to Synthetic Video. A Full Technical Report", technical report, arXiv:2607.12231, 2026.
   URL: https://arxiv.org/abs/2607.12231
5. **N. Cudlenco**, M. Masala, and M. Leordeanu, „Authoring for Living Worlds: Tool-Constrained LLM Agents for Executable Multi-Actor Scenarios", in European Conference on Computer Vision Workshops (ECCVW 2026), AIW: Agent in World.
   URL: https://openreview.net/forum?id=kFK4sVgN9b, https://arxiv.org/abs/2604.10383
6. **N. Cudlenco**, M. Masala, and M. Leordeanu, „GTASA: Ground Truth Annotations for Spatiotemporal Analysis, Evaluation and Training of Video Models", accepted for publication at the Asian Conference on Computer Vision (ACCV 2026).
   URL: https://arxiv.org/abs/2604.10385
7. **N. Cudlenco**, C. Orellano, M. Abushosha, M. Said, M. Plans Pujolras, and L. Ramirez-Lopez, „Self-Supervised Pretraining for Near-Infrared Chemical Property Prediction across Instruments", in Asian Conference on Computer Vision Workshops (ACCVW 2026), MM4SciReal: Multimodal AI for Science and the Real World.
   URL: https://openreview.net/forum?id=hiqKESmhqY



SOFTWARE AND DATA

- **GEST-Engine, open-source simulation engine** that validates, grounds and executes GESTs, with the procedural and agentic GEST generation and the production orchestration: github.com/ncudlenco/mta-sim, github.com/ncudlenco/multiagent_story_system
- **GTASA, multi-actor video corpus** with exact spatio-temporal ground truth: huggingface.co/datasets/nnc-001/gtasa-01
- **proximetricsR**, R package for spectral preprocessing and chemometric calibration of NIR sensors, CRAN, 2026 (co-author): CRAN.R-project.org/package=proximetricsR

SUMMER SCHOOLS

- SSIMA 2018, International Summer School on Imaging with Medical Applications, Sibiu, Romania, July 2018. Poster presentation.
- EEML 2019, Eastern European Machine Learning Summer School, Bucharest, Romania, 1-6 July 2019. Poster presentation.

RESEARCH PROJECTS

- „Spacetime Vision: Towards Unsupervised Learning in the 4D World", EEA and Norway Grant, code EEA-RO-2018-0496, 2019-2022, coordinated by Prof. Dr. Marius Leordeanu.
- „Vision in Words: Automatic Linguistic Description of Objects, People and their Interactions in Indoor Videos", UEFISCDI, code TE-2016-2182, 2018-2020, coordinated by Prof. Dr. Marius Leordeanu.

AWARDS

- UEFISCDI award, PRECISI-2020 competition (research results, articles), registration no. PN-III-P1-1.1-PRECISI-2020-52093, for the article „Reading into the mind's eye: Boosting automatic visual recognition with EEG signals" (Neurocomputing, 2020), December 2020.
- Member of the awarded research team, Romanian Research Gala 2024 (Gala Cercetării Românești), Ministry of Research, Innovation and Digitization, application no. 720062, February 2024.

RESEARCH AND PROFESSIONAL EXPERIENCE

Institute of Mathematics of the Romanian Academy — Bucharest, Romania

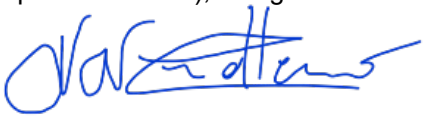
Research Assistant — Apr 2020 to present

- Built from scratch the GEST-Engine, about 100K lines of Lua, C++ and Python developed over six years. It takes a GEST as input, validates and parses it against the affordances of a 3D simulation engine, grounds its actors, objects and actions in the virtual world, orchestrates them under the graph's temporal constraints, and executes the result into synchronized multi-actor videos, recording exact 3D entity and camera state, pairwise spatial relations, instance segmentation, event-to-frame mappings and natural-language descriptions at every frame; it has since been extended with depth and actor pose. It runs on GTA San Andreas and is designed to be **engine-agnostic** (ICLR 2026 Workshop on World Models, full technical report on arXiv, ACCV 2026).
- Built the production orchestration that renders corpora unattended and **scales with the number of parallel virtual machines** (more than 10^{26} scene configurations before any action is chosen), and with it **the GTASA corpus**, a 938-scenario sample of what it can generate, with 892M pairwise spatial relations and 1.45M annotated frames, about **84 times more spatial relations per annotated frame than Action Genome**, the densest prior video dataset we found, released together with the engine (ACCV 2026).
- Built an **agentic system that goes from natural language to video with exact ground truth**: LLM agents turn a free-text story into a GEST, a semantic representation of who does what, with whom, where and when, building it only through more than 30 tools that validate every operation against the simulator's state. It executed 20 of 25 stories end to end where a staged LLM pipeline produced none of 50, and its graphs recover the source scenarios at 0.83 event F1 (ECCV 2026 Workshop AIW).
- Applied the engine and corpus to video understanding and video generation: human studies against frontier video generators, video captioning with synthetic pretraining, and a spatio-temporal probing benchmark of frozen video encoders, in which **relations between entities are decoded only slightly above chance** by all six encoders tested.

BÜCHI Labortechnik AG — Flawil, Switzerland

Lead ML Architect — Jan 2026 to present

- Set the **technical direction and engineering standards** of the Data Science team and of the company's NIR spectroscopy AI platform, owning the cloud and MLOps architecture and ML system design end to end, aligned with the software and hardware/firmware system architects and with product and R&D roadmaps.
- Architect the next generation of web services powering **NIR modelling** company-wide: embedded in all products, as a standalone web service, and as internal tooling for the application specialists who build customer models.
- Drive applied research and AI innovation, from **self-supervised encoders for NIR spectra**, pretrained on about 7.2 million unlabeled food and feed spectra (ACCV 2026 Workshop MM4SciReal), to agentic LLM systems, taking results from prototype to production.



- Mentor engineers in a five-person ML team, screen and interview candidates, and raise engineering standards through architecture and code reviews while staying hands-on in the codebase.

Machine Learning Engineer | MLOps Feb 2024 to Dec 2025

- Introduced and defined the MLOps role at the company, together with its least-privilege access model, and owned the team's full ML lifecycle, from data and cloud infrastructure through model packaging and deployment to monitoring.
- Took ownership of the system architecture of a new cloud data-science platform and aligned it with security best practices, validating the design with the software architecture lead and **cutting** cloud running costs by **about 30,000 CHF per year**.
- Built the cloud foundation on the company's Azure estate with Infrastructure as Code (Terraform, ARM templates) and Docker, and CI/CD pipelines with GitHub Actions that cut service deployment **from about 2 hours of manual work to under 5 minutes**.
- Delivered production service components and an internal portal consolidating all platform services. Defined environment, versioning, release and branching strategies, and ensured compliance with ISO 27001, ISO 31000 and GDPR.

Professional Software Engineer Jun 2023 to May 2024

- Developed software for laboratory instruments across the spectroscopy and chromatography domains and built and maintained CI/CD pipelines for these products.

Earlier industry experience

- **Acodis AG**, Winterthur, Switzerland, Experienced Software Engineer, Aug 2021 to Jun 2023: built in a proprietary purely functional language the back-end of a web application for automated data extraction using **intelligent document processing**, integrating and evaluating ML solutions from R&D.
- **Self-employed software consultant**, Bucharest, Jan 2018 to Aug 2021: designed and delivered software for infrastructure, construction and energy companies, including SNCF, Bouygues, EGIS and ELIA: plugins for the CAD and document-management tools their engineers use (AutoCAD, Inventor, ProjectWise, Microsoft Office), and desktop, web and cross-platform applications (C#, C++, WPF, Vue.js, .NET Web API, Java).
- **Yourshore and Graitec**, Bucharest, Software Engineer / Junior Software Engineer, Jul 2014 to Jan 2018: developed plugins for Autodesk Vault, Bentley ProjectWise and Microsoft Office, and built an optimized geometry library used by the development team to create 3D modeling tools for Autodesk Advance Steel (C#, C++, WPF).

TEACHING

- **Teaching Assistant, Analysis of Algorithms**, University Politehnica of Bucharest, Faculty of Automatic Control and Computer Science, Oct 2014 to Jan 2015 (course led by Dr. Matei Popovici): laboratory sessions and laboratory material, including a lab on Turing machines relating a Java implementation to the theoretical model.

SKILLS

Research areas: computer vision, vision and language, video understanding and generation, synthetic data, near-infrared spectroscopy and chemometrics, EEG-based visual recognition, self-supervised learning, LLM agents

Frameworks and tools: PyTorch, TensorFlow, Keras, Hugging Face, Docker, Azure, Terraform, GitHub Actions

Programming: Python, C#, C/C++, Java, JavaScript, R, Julia, functional languages

Languages: Romanian (native), English (fluent, C2), German (intermediate, B1), French and Russian (passive)

