Dmytro Kuzmenko

dmytro.o.kuzmenko@gmail.com
dmytro-kuzmenko.github.io
Google Scholar | GitHub

PhD Candidate in Computer Science | Research Engineer
Turin, Italy

PhD Candidate in Computer Science at the National University of Kyiv-Mohyla Academy (NaUKMA) and Visiting PhD Researcher at AIRlab, University of Turin. Dmytro works at the intersection of reinforcement learning, robot learning, and human-robot interaction, with emphasis on efficient and reliable embodied intelligence: model-based RL, policy distillation, modular multi-task robot-learning architectures, and socially aligned, human-aware robot behaviour. He combines research with production-oriented ML engineering experience.

Research Interests

Reinforcement learning for robotic control | model-based RL and policy distillation | modular / mixture-of-experts robot-learning architectures | vision-language-action (VLA) models | efficient embodied intelligence | human-robot interaction and safety-aware, ethically-grounded robot decision-making | semi-supervised visual perception for assistive and wearable robotics.

Education

PhD in Computer Science *2023 – present*
National University of Kyiv-Mohyla Academy (NaUKMA), Kyiv, Ukraine. Advisor: Dr. Nadiya Shvai. Expected defence: September 2026. Dissertation: *Resource-Aware Embodied Intelligence Through the Optimisation of Data, Compute, and Trust.*

MSc in Applied Mathematics *2021 – 2023*
National University of Kyiv-Mohyla Academy (NaUKMA), Kyiv, Ukraine.

BSc in Software Engineering *2016 – 2021*
National University of Kyiv-Mohyla Academy (NaUKMA), Kyiv, Ukraine.

Academic & Research Appointments

AIRlab, Department of Computer Science, University of Turin Turin, Italy
Visiting PhD Researcher Feb 2026 – present
Host: Prof. Cristina Gena. Research on assistive and human-aware robotics, human-robot teaming, and HRI evaluation. Contributor to *This/That: Spatial Reference Resolution for VLMs* (ongoing) and advisor on a completed scene-mapping project for a museum tour-guide robot.

Department of Multimedia Systems, NaUKMA Kyiv, Ukraine (Remote)
Senior Lecturer Oct 2022 – present
Designed and delivered seven graduate and undergraduate courses in machine learning, computer vision, reinforcement learning, and NLP (21 offerings to date); supervised 13 undergraduate theses and 25+ student research projects. See Teaching Experience and Supervision.

Neural Robotics Lab, University of Toronto Toronto, Canada (Remote)
Machine Learning Researcher Apr 2022 – Oct 2023
Collaborated on efficient visual perception for robotic exoskeleton and wearable locomotion systems; contributed semi-supervised stair recognition pipelines for human-robot walking environments. Work published at IROS 2023 and BioMedical Engineering Online 2024.

Industry Experience

Head of AI, OnlyMonster (Remote) *Jan 2025 – present*
Founding AI engineer; technical direction for core AI products, production-scale learning-based systems, and low-latency inference / moderation pipelines.

Data Science Engineering Manager, Technical Lead, Litslink (Kyiv, Ukraine) *Nov 2022 – Dec 2024*
Led applied ML projects across computer vision, predictive modelling, and LLM pipelines (25+ production and PoC projects, including Sportsbox AI).

Publications

Peer-Reviewed Journal Articles

[J1] Kuzmenko, D., Shvai, N. "MoIRA: Modular Instruction Routing Architecture for Multi-Task Robotics." *Neurocomputing*, vol. 674, 132962, 2026. doi:10.1016/j.neucom.2026.132962 | arXiv:2507.01843

[J2] Kurbis, A. G., Kuzmenko, D., Ivanyuk-Skulskiy, B., Mihailidis, A., Laschowski, B. "StairNet: Visual

Recognition of Stairs for Human-Robot Locomotion." *BioMedical Engineering OnLine*, 23(1):20, 2024. doi:10.1186/s12938-024-01216-0

[J3] Beimuk, V., Kuzmenko, D. "Energy Conservation for Autonomous Agents Using Reinforcement Learning." *NaUKMA Research Papers: Computer Science*, vol. 8, pp. 68–75, 2025. doi:10.18523/2617-3808.2025.8.68-75

Peer-Reviewed Conference Papers

[C1] Kuzmenko, D., Shvai, N. "When Robots Say No: The Empathic Ethical Disobedience Benchmark." In *Proc. ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Edinburgh, Scotland, pp. 168–176, March 2026. doi:10.1145/3757279.3785547 | arXiv:2512.18474 | code | project page

[C2] Kuzmenko, D., Shvai, N. "Knowledge Transfer in Model-Based RL Agents for Efficient Multi-Task Learning." In *Proc. 24th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Detroit, MI, USA, pp. 2597–2599, May 2025. doi:10.5555/3709347.3743949 | arXiv:2501.05329

[C3] Kuzmenko, D., Shvai, N. "Balancing Performance and Efficiency in Zero-shot Robotic Navigation." In *Proc. 19th International Conference on ICT in Education, Research and Industrial Applications (ICTERI 2024)*, Springer CCIS, 2025. doi:10.1007/978-3-031-81372-6_28 | arXiv:2406.03015

[C4] Kuzmenko, D., Tsepa, O., Kurbis, A. G., Mihailidis, A., Laschowski, B. "Efficient Visual Perception of Human-Robot Walking Environments using Semi-Supervised Learning." In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Detroit, MI, USA, October 2023. doi:10.1109/IROS55552.2023.10341654

Preprints

[P1] Kuzmenko, D., Shvai, N. "TD-MPC-Opt: Distilling Model-Based Multi-Task Reinforcement Learning Agents." arXiv preprint arXiv:2507.01823, 2025 (under review). arxiv.org/abs/2507.01823 | code

Workshop Papers & Posters

[W1] Kuzmenko, D., et al. "Semi-Supervised Learning for Efficient Visual Perception of Human-Robot Walking Environments." *Toronto Robotics Conference (TRC)*, Toronto, Canada, July 2024. researchgate.net

[W2] Kuzmenko, D., et al. "SSL for Efficient Perception of Human-Robot Walking Environments." *International Conference on Advanced Intelligent Robotics (ICAIR)*, 2023. pdf

[W3] Kuzmenko, D., et al. "SSL for Efficient Perception of Human-Robot Walking Environments." *ICRA Workshop on Computer Vision in Robotics Research (ICRA CVWR)*, London, UK, May 2023. pdf

Full and up-to-date publication list, BibTeX entries, and citation metrics available on Google Scholar *and* personal website.

Dissertation Research

The dissertation investigates efficient learning-based systems for robotic control and human-robot interaction, organised around four interconnected research threads:

Policy Distillation for Model-Based Multi-Task RL. Developed TD-MPC-Opt, a distillation framework transferring knowledge from a multi-task TD-MPC2 teacher to efficient student agents. Studies compute-efficient policy transfer, task generalisation under constrained inference budgets, and the structure of shared latent world-model representations across task families. [P1, C2]

Modular Instruction Routing for Multi-Task Robot Learning. Designed MoIRA, a mixture-of-experts routing architecture for task-conditioned multi-task robot policy learning. MoIRA learns to route task instructions to specialised policy sub-networks, improving modularity and per-task performance over monolithic baselines while enabling compositional task inference at test time. [J1]

Human-Aware Robot Behaviour and Ethical Disobedience. Constructed EED-Gym, a simulation benchmark and evaluation suite for empathic ethical disobedience in human-robot interaction – the conditions under which a robot should override or refuse a human command for safety or ethical reasons. Introduced empathy-aware refusal metrics and a dedicated HRI evaluation protocol. [C1]

Efficient Perception for Assistive Locomotion. Developed semi-supervised ViT-based visual perception pipelines for robotic exoskeleton and wearable locomotion systems, achieving 98.8% stair recognition accuracy with a ∼35% reduction in labelled data requirements. Established shared evaluation protocols for semi-supervised visual perception in human-robot walking environments. [J2, C4, W1–W3]

Other Research Projects & Activities

- **This/That: Spatial Reference Resolution for VLMs** (AIRlab, UniTo; contributor, ongoing) – studies how well vision-language models and 3D pipelines resolve deictic references ("this"/"that") from a person's point of view.
- **Museum Tour-Guide Robot – Scene Mapping** (AIRlab, UniTo; advisor, completed) – 3D scene mapping pipeline for an autonomous museum tour-guide robot.
- **Reachy Mini educational deployment** (AIRlab, UniTo) – introduced a Reachy Mini robot to the lab; the platform has since been used in a labmate's dissertation work on empathy, seeded a new Python course for non-technical master's students, and is used by a supervised thesis student to teach English to children.

Teaching Experience

All courses at the Department of Multimedia Systems, National University of Kyiv-Mohyla Academy (NaUKMA), Kyiv, Ukraine. Role in all offerings: Instructor (sole lecturer; course design, lectures, assessments, and lab supervision).

Reinforcement Learning *MSc (Graduate)*
Instructor – course design, lectures, assessments, lab supervision Spring 2023, 2024, 2025 (3 offerings)

Machine Learning *MSc (Graduate)*
Instructor – course design, lectures, assessments, lab supervision Fall 2022, Fall 2023, Spring 2025, Spring 2026 (4 offerings)

Image Analysis & Computer Vision *MSc (Graduate)*
Instructor – course design, lectures, assessments, lab supervision Spring 2024, 2025, 2026 (3 offerings)

NLP Systems *MSc (Graduate)*
Instructor – course design, lectures, assessments, lab supervision Fall 2024, Fall 2025 (2 offerings)

Computer Vision *MSc (Graduate)*
Instructor – course design, lectures, assessments, lab supervision Fall 2022, 2023, 2024, 2025 (4 offerings)

Machine Learning *BSc (Undergraduate)*
Instructor – course design, lectures, assessments, lab supervision Fall 2024, Fall 2025 (2 offerings)

Deep Learning for Computer Vision *BSc (Undergraduate)*
Instructor – course design, lectures, assessments, lab supervision Spring 2024, 2025, 2026 (3 offerings)

Total: 7 distinct courses across 21 offerings (October 2022 – present).

Supervision

- **Undergraduate Theses.** Supervised 13 undergraduate theses at NaUKMA (2022 – present), including student-led work resulting in co-authored publications, e.g. Beimuk, V., Kuzmenko, D., "Energy Conservation for Autonomous Agents Using Reinforcement Learning," *NaUKMA Research Papers: Computer Science*, 2025 [J3]; Severhin, O., Kuzmenko, D., Shvai, N., "Efficient Policy Learning via Knowledge Distillation for Robotic Manipulation," NaUKMA, 2025; Voitishyn, M., Kuzmenko, D., "A Hybrid AI Model for Financial Market Prediction," NaUKMA, 2025.
- **Undergraduate & Postgraduate Research Projects.** Supervised 25+ student research projects at NaUKMA (2022 – present), covering reinforcement learning, computer vision, NLP, and robotics perception; several projects submitted to conference submissions.

Academic Service

- Program Committee Member, *International Conference on ICT in Education, Research and Industrial Applications (ICTERI 2026)*
- Reviewer, *International Conference on Social Robotics (ICSR 2026)*
- Reviewer, *ACM/IEEE International Conference on Human-Robot Interaction (HRI 2026)*
- Reviewer, *Journal of Field Robotics* (Wiley), 2025 – 2026

Invited Talks & Conference Presentations

- **When Robots Say No: The Empathic Ethical Disobedience Benchmark.** Presentation, *COMETE Workshop*, University of Turin. Apr 2026
- **When Robots Say No: The Empathic Ethical Disobedience Benchmark.** Main-track paper presentation, *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Edinburgh, Scotland. Mar 2026
- **Robotics and AI research.** Talk and panel participation, *Kyiv AI & BigData Day 2025.* Sep 2025
- **Knowledge Transfer in Model-Based RL Agents.** Invited talk, RL Reading Club, *University of Southampton*, UK. Jun 2025
- **Knowledge Transfer in Model-Based RL Agents for Efficient Multi-Task Learning.** Poster, *24th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Detroit, MI, USA. May 2025
- **SSL for Efficient Visual Perception of Human-Robot Walking Environments.** Research presentation, *Toronto Robotics Conference (TRC)*, Toronto, Canada. Jul 2024
- **SSL for Efficient Perception of Human-Robot Walking Environments.** Workshop paper presentation, *ICRA Workshop on Computer Vision in Robotics Research (CVWR)*, London, UK. May 2023
- **SSL for Efficient Perception of Human-Robot Walking Environments.** Conference paper presentation, *International Conference on Advanced Intelligent Robotics (ICAIR).* 2023

Research Visits & Schools

- **Visiting PhD Researcher**, AIRlab, Department of Computer Science, University of Turin, Turin, Italy. Feb 2026 – present
- **PhD Summer School in Artificial Intelligence and Creativity.** University of Bologna, residential school, Bertinoro, Italy. Jul 2024

Outreach & Professional Service

- Micro-Committee Member, AI HOUSE Oct 2024 – Sep 2025
- Data Science Mentor, Projector Institute May 2022 – Aug 2023
- Machine Learning Mentor, Women Who Code Kyiv Aug 2022 – Feb 2023
- Machine Learning Lead, Faculty of Computer Science, NaUKMA Feb 2022 – Jul 2022
- Contributed a module on AI in cybersecurity for Google's "AI for Productivity" programme
- Delivered over 20 talks at AI/ML meetups and conferences in Ukraine and abroad

Technical Skills

Research / Robotics: model-based RL, policy distillation, imitation learning, VLA/VLM systems, embodied AI benchmarking, HRI evaluation, computer vision for robotics
Simulation / Robotics Tools: MuJoCo, Gymnasium, Isaac Gym/Lab, Genesis, PyBullet, Habitat, ROS/ROS2
Programming / ML: Python, PyTorch, HuggingFace, TensorFlow, Docker, Git
Systems: vLLM, low-latency inference, cloud deployment (Azure / AWS / GCP), production ML systems

Languages

English: Professional proficiency (C1)
German: Independent user (B2)
Ukrainian: Native

References

Available upon request.