Alexander Pugantsov

Padua, Veneto, Italy | a.pugantsov@icloud.com | pugantsov

Machine Learning Researcher | Trustworthy AI, Fairness, Privacy, Evaluation
Research-oriented Applied ML Scientist with experience building rigorous evaluation frameworks for machine learning systems under deployment constraints, with particular focus on fairness and privacy in attribute-unaware settings, robustness under shift, and uncertainty-aware assessment. Combines academic research depth with applied experience in retrieval, NLP, and LLM-enabled knowledge systems.

EXPERIENCE

University of Padua Oct. 2024 – Present
Postdoctoral Researcher Padua, Veneto, Italy

- Built reproducible evaluation pipelines for auditing fairness and group-level outcomes when protected demographic attributes are unavailable or restricted, reflecting realistic deployment constraints.
- Designed and ran large-scale empirical studies across classification, ranking, and vision settings to evaluate robustness under distribution shift, uncertainty, and limited-label conditions.
- Analyzed privacy risks in aggregate demographic estimation by developing adversarial query strategies and studying privacy-utility trade-offs under differential privacy-based defenses.

Confidential Industry Client (NDA) Apr. 2023 – Sept. 2024
Machine Learning Engineer Remote

- Designed and implemented a semantic retrieval pipeline combining dense retrieval and cross-encoder reranking, improving search quality and supporting deployment in a production workflow.
- Constructed a domain ontology from noisy hierarchical data using LLM-assisted enrichment and graph-based clustering, improving knowledge organization and downstream navigation.
- Implemented domain-specific acronym disambiguation using PEFT/LoRA with quantized LLMs to support high-precision query expansion in acronym-heavy corpora.

EDUCATION

University of Glasgow Dec. 2019 – Jun. 2024
Ph.D., Computing Science Glasgow, United Kingdom

Thesis: *Effective Intermediate Task Selection in Transfer Learning*

- Developed methods for predicting effective task sequences in transfer learning without exhaustive retraining, reducing training cost by up to 96.5% while preserving downstream performance.

University of Glasgow Sep. 2014 – Jun. 2019
M.Sci., Software Engineering Glasgow, United Kingdom

Dissertation: *Deep Learning & User Modeling*

SELECTED PUBLICATIONS & ONGOING RESEARCH

- A. Pugantsov, A. Fabris, G.A. Susto, A. Moreo. Manuscript on privacy-preserving fairness auditing and quantification-based evaluation. *Under review*. 2026.
- A. Pugantsov, R. McCreadie: Divergence-Based Domain Transferability for Zero-Shot Classification. *Findings of the Association for Computational Linguistics (EACL)*. 2023.
- A. Pugantsov, R. McCreadie: Identifying Suitable Tasks for Inductive Transfer Through the Analysis of Feature Attributions. *European Conference on Information Retrieval (ECIR)*. 2022.

SKILLS

- **Programming / ML**: Python, PyTorch, scikit-learn, Hugging Face
- **Areas**: Trustworthy AI, fairness, privacy-preserving ML, uncertainty, robustness under distribution shift, retrieval/ranking, NLP, representation learning
- **Tools**: SQL, Docker, Linux, Git, Pandas, NumPy
- **Languages**: English, Russian, Italian