

# Alexander Pugantsov

## RESEARCH AREAS

---

- Quantification learning and class-prevalence estimation.
- Machine learning under dataset and distribution shift.
- Evaluation and robustness under changing data distributions.
- Transfer learning, domain transferability, and task similarity.
- Population-level inference from unlabeled data.
- Privacy, security, and fairness applications of quantification.

## RESEARCH PROFILE

---

Machine learning researcher specializing in quantification learning, dataset shift, transfer learning, and the evaluation of machine learning systems under changing data distributions. My research investigates how models and population-level estimators behave when training and deployment distributions differ, including methods for estimating class prevalence in unlabeled populations, evaluating robustness under distribution shift, and understanding the reliability and security of aggregate predictions.

From 2024 to 2026, I conducted postdoctoral research within the “QuaDaSh – Quantification in the Context of Dataset Shift” project at the University of Padua, developing and evaluating quantification-based approaches under realistic distributional change. This work includes methodological and empirical research on quantification, evaluation protocols, robustness, population-level inference, and the privacy and security implications of repeated aggregate estimation.

My earlier doctoral research at the University of Glasgow investigated transferability across tasks and domains, including methods for predicting when knowledge learned from one data distribution can be successfully transferred to another. This included research on divergence-based measures of domain transferability, zero-shot classification, task similarity, and representation-based transfer learning.

Together, this work forms a continuous research programme concerned with learning, estimation, and evaluation when the distribution encountered at deployment differs from that observed during training.

### Academic Profiles and Identifiers

- **Scholar:** [https://scholar.google.com/citations?user=b0i\\_tFYAAAAJ&hl=en](https://scholar.google.com/citations?user=b0i_tFYAAAAJ&hl=en)
- **Scopus:** <https://www.scopus.com/authid/detail.uri?authorId=58114630300>
- **ORCID:** <https://orcid.org/0000-0001-6630-0258>

## ACADEMIC EMPLOYMENT

---

- **Postdoctoral Researcher, University of Padua**

Padua, Italy  
Oct 2024 – May 2026

*Quantification and Machine Learning Under Dataset Shift*

Advisor: Prof. Gian Antonio Susto.

  - Conducted postdoctoral research within the “QuaDaSh – Quantification in the Context of Dataset Shift” PRIN-PNRR project, funded by NextGenerationEU.
  - Developed and evaluated quantification methods for estimating class prevalence in unlabeled populations under distributional change.
  - Designed experimental methodologies and evaluation protocols for studying quantification performance under dataset shift and deployment constraints.
  - Investigated robustness, reliability, and failure modes of population-level estimators under changing data distributions.
  - Developed applications of quantification to fairness and diversity assessment when individual-level sensitive attributes are unavailable.
  - Investigated privacy and security risks arising from repeated access to aggregate quantification outputs, together with privacy-preserving defensive mechanisms.

## EDUCATION

---

- **Ph.D. in Computing Science, University of Glasgow** Glasgow, United Kingdom  
*A Framework for Effective Intermediate Task Selection in Transfer Learning* Dec 2019 – June 2024  
Advisors: Dr. Richard McCreadie, Dr. Jeff Dalton.  
**Final award:** Doctor of Philosophy, awarded 25 June 2024.  
**Assessment:** Research Ph.D. degrees are not assigned a numerical grade or degree classification and are not credit-rated under the SCQF.
  - Research focused on transfer learning, domain transferability, task similarity, and methods for predicting whether knowledge learned from one task or data distribution would transfer effectively to another.
  - Developed divergence-based methods for measuring relationships between source and target domains and predicting transferability for zero-shot classification.
  - Investigated representation-based and feature-attribution-based measures of task similarity for sequential and inductive transfer learning.
  - Research addressed model behaviour under differences between source and target data distributions, providing methodological foundations for my later work on machine learning under dataset shift.
- **M.Sci. in Software Engineering, University of Glasgow** Glasgow, United Kingdom  
*Deep Learning & User Modelling* Sep 2014 – June 2019  
Advisor: Prof. Joemon Jose.  
**Final classification:** Upper Second-Class Honours (2:1; officially “with Honours of the Second Class, Division i”).  
**Qualification level:** Five-year integrated Master’s degree, SCQF Level 11.  
**Credits:** 610 SCQF credits (305 ECTS).  
**Grading scale:** University of Glasgow 22-point scale (22 = highest, 0 = lowest). A = Excellent, B = Very Good, C = Good, D = Satisfactory. Within each letter grade, a lower numerical suffix indicates a higher result (e.g., A1 = 22/22, A5 = 18/22, B1 = 17/22, B3 = 15/22).
  - **Selected academic results:**
    - \* MSci Research Proposal and Project (Level 5 Dissertation) – A5 (18/22; 80 SCQF credits / 40 ECTS)
    - \* Individual Project (Level 4 Dissertation) – A4 (19/22; 40 SCQF credits / 20 ECTS)
    - \* Cyber Security Fundamentals – A4 (19/22; 10 SCQF credits / 5 ECTS)
    - \* IT Architecture – A3 (20/22; 10 SCQF credits / 5 ECTS)
    - \* Research Methods and Techniques – B2 (16/22; 10 SCQF credits / 5 ECTS)
    - \* Artificial Intelligence – B3 (15/22; 10 SCQF credits / 5 ECTS)
  - Conducted research on recommender systems and user modelling using large-scale financial social-media data, investigating deep neural architectures for learning user and content representations.
  - Developed and evaluated sequential LSTM- and CNN-based models and attention-based architectures for modelling user behaviour and information preferences.
  - Investigated enrichment of learned representations using external knowledge graphs, including Wikidata and DBpedia, queried using SPARQL.

## PUBLICATIONS AND RESEARCH OUTPUTS

---

### Journal Articles

- A. Pugantsov, A. Fabris, G.A. Susto, A. Moreo: PAIRE: A Testbed for Measuring Dataset Diversity and Model Fairness Under Limited Sensitive Information. *Information Processing & Management (IP&M)*. 2027 (published online 2026). DOI: [10.1016/j.ipm.2026.105095](https://doi.org/10.1016/j.ipm.2026.105095).  
**Journal metrics:** JCR 2025: Q1, JIF 8.1; SCImago 2025: SJR 1.845, Q1.

### Conference Papers

- A. Pugantsov, R. McCreadie: Divergence-Based Domain Transferability for Zero-Shot Classification. *Findings of the Association for Computational Linguistics (EACL)*. 2023. DOI: [10.18653/v1/2023.findings-eacl.122](https://doi.org/10.18653/v1/2023.findings-eacl.122).  
**Conference ranking:** CORE 2023: A.
- A. Pugantsov, R. McCreadie: Identifying Suitable Tasks for Inductive Transfer Through the Analysis of Feature Attributions. *European Conference on Information Retrieval (ECIR)*. 2022. DOI: [10.1007/978-3-030-99739-7\\_16](https://doi.org/10.1007/978-3-030-99739-7_16).  
**Conference ranking:** CORE 2023: A.

### Shared-Task and Evaluation Reports

- A. Pugantsov, R. McCreadie: University of Glasgow Terrier Team (uogTr) at the TREC 2021 Incident Streams Track. *Proceedings of the Thirtieth Text REtrieval Conference (TREC 2021)*, NIST Special Publication 500-335. 2021. DOI: [10.6028/NIST.SP.500-335.incident-uog\\_tr\\_team](https://doi.org/10.6028/NIST.SP.500-335.incident-uog_tr_team).
- A. Pugantsov, R. McCreadie: University of Glasgow Terrier Team (uogTr) at the TREC 2020 Incident Streams Track. *Proceedings of the Twenty-Ninth Text REtrieval Conference (TREC 2020)*, NIST Special Publication 1266. 2020. DOI: [10.6028/NIST.SP.1266.incident-UoGTr](https://doi.org/10.6028/NIST.SP.1266.incident-UoGTr).

### Manuscripts and Ongoing Research

- A. Pugantsov, A. Fabris, G.A. Susto, A. Moreo: Quantification as a Service: Privacy Risks in Aggregate Prevalence Estimation. Manuscript in preparation. 2026.

## RESEARCH PROJECTS

---

- **QuaDaSh: Quantification in the Context of Dataset Shift (2024–2026)**

PRIN-PNRR project funded by NextGenerationEU. Grant No.: P2022TB5JF; CUP: C53D23007970001.

My research within QuaDaSh focused on quantification methods for estimating class prevalence in unlabeled populations when training and deployment distributions differ. I developed experimental methodologies and evaluation protocols for studying quantification under dataset shift; investigated robustness and reliability under deployment conditions; developed applications to fairness and diversity assessment under attribute unavailability; and studied the privacy and security implications of repeated aggregate prevalence estimation.

- **Quantification for Attribute-Unaware Dataset Diversity and Model Fairness Evaluation (2024–2026)**

Developed quantification-based approaches for estimating otherwise unavailable population characteristics, including demographic prevalence, in support of fairness and diversity assessment. This work examines how population-level estimates can support evaluation when individual sensitive attributes are unavailable.

- **Privacy and Security of Quantification Systems (2025–2026)**

Investigated information leakage from repeated access to aggregate class-prevalence estimates returned by deployed quantification services. Developed attack and evaluation methodologies and studied privacy-preserving mechanisms for limiting leakage while retaining useful population-level estimates.

## RESEARCH SOFTWARE AND DATASETS

---

- **PAIRE – Dataset Diversity and Fairness Evaluation Testbed (2026)**

<https://github.com/pugantsov/PAIRE>

Open-source research testbed and reproducibility framework accompanying the PAIRE study, supporting quantification-based diversity estimation, fairness assessment, and adversarial vulnerability evaluation.

- **TREC Fair Ranking 2022 – Single-Region Subset for Quantification and Fairness Estimation (2026)**

DOI: [10.5281/zenodo.19476763](https://doi.org/10.5281/zenodo.19476763)

Published derived dataset retaining English Wikipedia articles assigned to exactly one subcontinental geographic region, for quantification and fairness evaluation with geographic region treated as the sensitive attribute.

- **PAIRE – UCI Adult Experimental Indices (2026)**

DOI: [10.5281/zenodo.14283869](https://doi.org/10.5281/zenodo.14283869)

Published reproducibility artifact containing the dataset indices used for the UCI Adult experiments in the PAIRE study.

- **Multi-Domain Yelp Business Reviews (2022)**

<https://github.com/pugantsov/multi-domain-yelp-business-reviews>

Multi-domain sentiment-classification dataset and construction toolkit derived from the Yelp Dataset Challenge, which contains text reviews and ratings for businesses across multiple categories.

## CONFERENCE PRESENTATIONS AND PARTICIPATION

---

- **17th Conference of the European Chapter of the Association for Computational Linguistics (EACL) (2023)**

Presented “Divergence-Based Domain Transferability for Zero-Shot Classification” as a poster, with an accompanying recorded research presentation published online. DOI: [10.48448/mz42-gy60](https://doi.org/10.48448/mz42-gy60).

- **44th European Conference on Information Retrieval (ECIR) (2022)**

Presented “Identifying Suitable Tasks for Inductive Transfer Through the Analysis of Feature Attributions” as a poster.

## TEACHING EXPERIENCE

---

- **Teaching Assistant, University of Glasgow**

Glasgow, United Kingdom

*School of Computing Science*

*Oct 2017 – May 2024*

- Provided paid teaching, tutorial, laboratory, assessment, and student-support activities across undergraduate and postgraduate Computing Science courses over seven academic years.
- Responsibilities included tutorials and laboratory sessions, coursework assessment and grading, examination marking at Honours and Master’s level, formative feedback, and technical and conceptual support for students.
- **Data Science and Analytics:** Big Data; Information Retrieval; Information Visualisation; Introduction to Data Science with Python.
- **Programming and Software Development:** Computing Fundamentals; Introductory Programming in Python; Software Engineering.
- **Cybersecurity:** Enterprise Cyber Security; Cyber Systems Forensics.
- **Postgraduate Teaching and Project Support:** Human–Computer Interaction; MSc Summer Projects. Provided scheduled drop-in project support to Master’s students, advising across multiple dissertation projects on research questions, hypotheses, methodology, implementation, and dissertation development.

## ACADEMIC SERVICE AND LEADERSHIP

---

- **Peer Reviewer (Ad hoc), Information Processing & Management (IP&M)** (2026)
- **Founder/Organizer, Machine Learning in Science Colloquium** (2020–2023)

Founded and organized a multidisciplinary seminar series on machine learning topics attended by staff and research students from multiple universities, culminating in a student-run conference featuring speakers from DeepMind, the Max Planck Institute, and the University of Edinburgh.

## RESEARCH GROUPS

---

- **Artificial Intelligence, Machine Learning & Control** (2024–2026)

University of Padua research group led by Prof. Gian Antonio Susto. My work in the group focused on quantification learning, machine learning under dataset shift, population-level estimation, robustness and evaluation, and applications to trustworthy and responsible machine learning.

- **Glasgow Information Retrieval Group (Terrier Team)** (2019–2024)

Research group at the University of Glasgow working on information retrieval, machine learning, and natural language processing. My doctoral research focused on transfer learning, domain transferability, representation learning, task similarity estimation, and predicting transfer performance across source and target domains.

## NON-ACADEMIC EMPLOYMENT

---

### Machine Learning Engineer, NDA

Remote

• *Applied machine learning, semantic retrieval, and language models.*

*Apr 2023 – Sep 2024*

- Developed and evaluated machine-learning methods for semantic information retrieval and knowledge organization.
- Improved semantic search pipelines using dense/contrastive representations and cross-encoder reranking for more accurate retrieval.
- Developed LLM-assisted approaches for constructing and enriching structured ontologies from noisy hierarchical data, combining language-model outputs with graph-based clustering and topic-modelling methods.
- Developed methods for acronym and terminology disambiguation using parameter-efficient fine-tuning (PEFT/LoRA) of quantized large language models, supporting structured query expansion and retrieval.
- Conducted end-to-end machine-learning development spanning problem formulation, experimentation, model evaluation, implementation, and integration into production-oriented engineering workflows.

### Automation Engineer, Barclays

Glasgow, United Kingdom

• *Compliance workflow automation and internal tooling.*

*Aug 2019 – Dec 2019*

- Developed internal software and automation tooling supporting screening and compliance workflows within a large financial-services environment.
- Designed workflow automation aimed at reducing manual processing and improving the integration of internal compliance processes.

### Software Engineering Intern, Scoop Markets

Glasgow, United Kingdom

• *Financial analytics and reporting workflow integration.*

*Jun 2018 – Sep 2018*

- Developed software supporting financial-data analytics and recommendation workflows.
- Built Go-based aggregation tooling for integrating and processing outputs produced by recommender-system components.

## LANGUAGES

---

- **English:** Native proficiency.
- **Russian:** Native proficiency.
- **Italian:** Intermediate working proficiency; reading comprehension and everyday/professional communication.