

Mateusz Krubinski, PhD

mkrubinski@gmail.com

+48 604 302 964

<https://www.linkedin.com/in/mateusz-krubi%C5%84ski-153a46198>

<https://scholar.google.com/citations?user=12NoTUUAAAAJ>

EXPERIENCE

Snowflake

Warsaw, Poland

July 2024 - October 2025 | Research Scientist

Worked on Cortex AISQL, Snowflake's multimodal platform providing language, vision, and speech intelligence.

- **Machine Translation:** Led the development of a next-generation translation system, directing a team of two engineers. Drove the technical strategy, designed comprehensive evaluation suites, implemented data collection and filtering pipelines, and orchestrated training and fine-tuning efforts.
- **Multilingual Entity Sentiment:** Led the end-to-end development and production rollout of a multilingual entity sentiment model, managing the entire project lifecycle — from product requirements and data generation to legal review, human evaluation, and deployment.
- **Audio Transcription:** Contributed to the initial launch of the transcription system, leading quality evaluation and implementing core audio processing components, including voice activity detection and audio chunking.
- **Multimodal LLM Enablement:** Core engineering work enabling multimodal LLM inference over visual data in a previously text-only platform. Implemented support for multi-image inputs.

Technologies: Python, PyTorch, Kubernetes, Go, vLLM, Hugging Face, Snowflake

Amazon Development Center

Bucharest, Romania

July 2022 - November 2022 | Applied Scientist Intern

Worked on automated processing of document images and scans. The internship led to a research paper presented at the AAAI 2024 Workshop: "Watermark Text Pattern Spotting in Document Images."

- Developed a custom architecture combining Faster R-CNN with a Transformer decoder.
- Built a full pipeline for synthetic data generation and designed protocols for manual annotation and automated evaluation.

Technologies: Python, PyTorch, AWS, Detectron2

Samsung R&D Institute Poland

Warsaw, Poland

April 2019 - September 2020 | Junior NLP Engineer

July 2018 - April 2019 | NLP Intern

Worked across the full Machine Translation pipeline – from research (pre-training, training, fine-tuning) to deployment (cloud serving, on-device compression).

- First author of a system description paper published at WMT 2020, having trained the top-performing MT models for Polish ↔ English directions.

Technologies: *Python, TensorFlow, Docker, Bash, Android SDK*

EDUCATION

Charles University

Prague, Czech Republic

October 2020 - September 2024 | PhD Student

- Worked on effective ways of fusing textual and visual features and evaluating the quality of multi-modal output.
- Grants:
 - CUNI GA UK (PI) - exploring the math capabilities of LLMs. The outcome of the project is a research paper presented at the NeurIPS 2023 workshop, i.e., “Basic Arithmetic Properties in the Space of Language Model Prompts”.
 - WELCOME (H2020) - text-to-text and speech-to-text translation of Arabic dialects. Co-organizer of the IWSLT 2024/2025 shared task on low-resource speech translation.

Warsaw University of Technology

Warsaw, Poland

October 2013 - July 2018 | Mathematics MSc

SELECTED PUBLICATIONS

Google Scholar (01.2026) - citations: 94, h-index: 5

- **Mateusz Krubinski** and Pavel Pecina, *Towards Unified Uni- and Multi-modal News Headline Generation*. In Findings of the Association for Computational Linguistics: EACL 2024, March 2024
- **Mateusz Krubinski** and Pavel Pecina, *MLASK: Multimodal Summarization of Video-based News Articles*. In Findings of the Association for Computational Linguistics: EACL 2023, May 2023
- **Mateusz Krubinski**, Erfan Ghadery et al., *Just Ask! Evaluating Machine Translation by Asking and Answering Questions*. In Proceedings of the Sixth Conference on Machine Translation, November 2021