

# Sebastian Loeschcke

+45 60 72 09 48 | [Email](#) | [LinkedIn](#) | [Scholar](#) | [GitHub](#) | [Website](#)

Ph.D. researcher in efficient ML with first-author publications at **NeurIPS**, **ICML**, and **ECCV**, and two Best Paper Awards. Research focuses on training efficiency and model compression through quantization and low-rank methods.

## EDUCATION

---

### University of Copenhagen

Copenhagen, Denmark

*Ph.D. in Computer Science*

*2023 – Present*

- Research on efficient training and inference for deep learning, with emphasis on low-rank methods and quantization. Advisors: Serge Belongie and Michael J. Kastoryano.
- Completed a six-month research visit in 2025 with Anima Anandkumar's group at Caltech as part of the Ph.D.

### Aarhus University

Denmark

*Bachelor's and Master's in Computer Science, Grade Averages: 11.1/12 (Bachelor's), 12/12 (Master's) 2018 – 2023*

## SELECTED FIRST-AUTHOR PUBLICATIONS

---

### TensorGRaD: Robust Tensor Decomposition for Memory-Efficient Neural Operator Training | *arXiv*, 2025

- **Sebastian Loeschcke**, David Pitt, Robert J. George, Jiawei Zhao, Cheng Luo, Yuandong Tian, Jean Kossaifi, Anima Anandkumar.
- Developed a robust low-rank and sparse tensor decomposition for memory-efficient neural operator training, reducing training memory by > 50% on neural operator workloads while matching Adam performance.

### LoQT: Low-Rank Adapters for Quantized Pretraining | *NeurIPS*, 2024

- **Sebastian Loeschcke**, Mads Toftrup, Michael J. Kastoryano, Serge Belongie, Vésteinn Snæbjarnarson.
- Proposed a quantized low-rank pretraining approach that matches full-precision Adam training in perplexity while reducing memory, enabling 13B training on < 24GB.

### Coarse-To-Fine Tensor Trains for Compact Visual Representations | *ICML*, 2024

- **Sebastian Loeschcke**, Dan Wang, Christian Leth-Espensen, Serge Belongie, Michael J. Kastoryano, Sagie Benaim.
- Developed a coarse-to-fine quantized tensor train method, a higher-order tensor decomposition for compact visual representations, improving optimization for image fitting, 3D fitting, and novel view synthesis.

### Text-Driven Stylization of Video Objects | *ECCV Workshop*, 2022/2023

- **Sebastian Loeschcke**, Serge Belongie, Sagie Benaim.
- Presented one of the first methods for text-based stylization of video objects, producing temporally consistent edits from user-specified text prompts.

### Progressive Parameter Space Visualization for Task-Driven SAX Configuration | *EuroVA*, 2020

- **Sebastian Loeschcke**, Marius Hognräfer, Hans-Jörg Schulz.
- Built a progressive visual analytics tool for hyperparameter search in SAX/PAA, helping users explore task-driven tradeoffs between compression and utility.

## EXPERIENCE

---

### Machine Learning Engineer Intern

Jul 2025 – Jan 2026

*Qualcomm AI Research*

*Amsterdam, Netherlands*

- Developed quantization-aware training methods using lattice vector quantization for large language models, training models up to 3B parameters on 5B tokens under the supervision of Markus Nagel.
- Reduced lattice vector quantization QAT overhead by  $\sim 2\times$  via distributed sharding and custom Triton kernels, achieving near integer-scalar QAT speed ( $\sim 5.0$  vs.  $4.3$  s/iter on  $8\times A100$ ).
- Showed that 2-bit weight lattice vector quantization can match full-precision perplexity, while substantially improving MMLU and common reasoning benchmarks through QAT.

### Machine Learning Engineer Intern

Jan 2023 – Aug 2023

*AMCS Group Vision AI*

*Copenhagen, Denmark*

- Developed and deployed real-time computer vision models for object detection, covering model training, cloud deployment, and production performance optimization in collaboration with cross-functional teams.

## Junior Data Scientist

May 2021 – Dec 2022

*Systematic*

*Aarhus, Denmark*

- Developed machine learning solutions for Systematic’s software products, including speech and gesture recognition models and anomaly detection methods for log data analysis.

## Teaching Assistant

2020 – 2026

*Aarhus University; University of Copenhagen*

*Denmark*

- Advised students on research projects in the pre-talent track at Aarhus University (2020–2021).
- Taught Vision and Image Processing at the University of Copenhagen (2025, 2026).

## Course Teacher, Data Analysis and Machine Learning

Jan – Jun 2026

*University of Copenhagen*

*Copenhagen, Denmark*

- Taught Data Analysis and Machine Learning in the Master’s programme in Informatics Education.

## AWARDS & TALKS

---

### Awards & Fellowships

- Best Paper Award (LoQT), WANT@ICML 2024.
- Best Paper Award (Text-Driven Stylization of Video Objects), CVEU@ECCV 2022.
- Awarded the DDSA PhD Fellowship, a fully funded Danish Data Science Academy program for doctoral researchers.
- Recipient of Queen Margrethe II’s travel grant from Aarhus University, 2022.

### Selected Talks

- Spotlight presentation of TensorGRaD at SAINTS 2026, Copenhagen.
- Invited talk, “Low-Rank and Quantized Learning,” at UC Santa Barbara, hosted by Zheng Zhang, 2025.
- Invited talk at NVIDIA Research on LoQT, 2024; oral presentations at WANT@ICML 2024 and CVEU@ECCV 2022.

### Academic Service

- Reviewer for **ICML**, **NeurIPS**, and **ICLR**.

## TECHNICAL SKILLS

---

**Core:** Quantization, low-rank methods, distributed training, model compression, efficient training and inference **Tools:** PyTorch, TensorFlow, TensorLy, Quimb **Programming:** Python, C/C++, SQL, Bash, Java **Systems:** Linux, Docker, Git