

Markus Hiller

Research Fellow in Artificial Intelligence, The University of Melbourne

markush.research@gmail.com | (+61) 0492 828326 | Melbourne, Australia | [Google Scholar](#)

PROFILE

Research Fellow in Artificial Intelligence working on efficient and low-resource deep learning for vision and language. Published contributions span efficient neural architectures, few-shot and low-data image classification, and multi-object tracking, and, through doctoral co-supervision, diffusion-based medical image generation and 3D reconstruction. Work has appeared at NeurIPS, TPAMI, ICLR, and ICML.

EDUCATION

Doctor of Philosophy, Computer Science (Artificial Intelligence)

04/2025

The University of Melbourne

Thesis: *Improving the Efficiency of Deep Visual Representation Learning*. Supervised by Prof. Tom Drummond and Dr. Krista A. Ehinger (University of Melbourne), and A/Prof. Mehrtaash Harandi (Monash University).

Focus on efficient deep visual representation learning, few-shot learning, and neural architecture design.

Master of Science, Mechatronics (with distinction)

05/2017

Friedrich-Alexander University, Germany

Specialisation in robotics, real-time systems, sensor data fusion, and estimation theory.

Thesis: *Global SLAM by Fusion of Laser, Odometry and Map Information*.

Bachelor of Science, Mechatronics

03/2015

Friedrich-Alexander University, Germany

Specialisation in control theory and sensor systems. Thesis: *Design of a Repetitive State-Space Controller for a Diesel Generator Based on the Internal Model Principle* (in German)

RESEARCH EXPERIENCE

Research Fellow in AI (Level B)

09/2024 – present

The University of Melbourne

- Investigating the integration of an external, searchable memory with generative models as a more efficient alternative to fully parametric approaches, spanning language and retrieval-guided image generation, with designs structured to transfer to multimodal and vision-language settings.
- Designing efficient architectures, including hybrid state-space and Transformer designs, and exploring geometry-inspired optimisation, continual learning, and distribution shifts across vision and language.
- Co-supervising doctoral research in diffusion-based medical image generation and reconstruction, as well as in efficient vision state-space models (see Teaching and Supervision).

Graduate Researcher (PhD Candidate)

04/2020 – 09/2024

The University of Melbourne

- Designed an efficient attention-based architecture that scales linearly with input sequence length and needs roughly 200x less compute than its closest competitor (Perceiver IO) for equal or better performance, generalising across images, point clouds, and text (NeurIPS 2024).
- Developed a few-shot image classification method that learns what matters most in the input at test time and reduces label bias from previously seen data, improving on prior work by up to 4.4% using self-supervised pretraining to avoid representation collapse (NeurIPS 2022).
- Improved new-task adaptation speed for gradient-based few-shot meta-learning by up to 61 percentage points without additional parameters, by encouraging a well-conditioned parameter space using local curvature and condition-number information (NeurIPS 2022).
- Co-developed a Transformer-based multi-object tracking method that removes hand-crafted post-processing by attending jointly to new observations and latent history over time (TPAMI 2022).

Wissenschaftlicher Mitarbeiter (Research & Teaching Associate)

05/2017 – 12/2019

Friedrich-Alexander University, Germany

- Led and coordinated environment-perception and navigation research within *FORobotics*, a joint project of 8 academic and 20 industry partners funded by the Bavarian Research Foundation (BayFOR), achieving the highest possible rating in all three annual funding-agency reviews.

- Researched semantic environment representations for intelligent navigation, yielding publications at robotics venues including IROS, ISR, and ECMR.
- Worked directly with industry partners, helping Bertrandt AG interpret and translate computer vision and SLAM research into their engineering projects.
- Co-led the institute's Robotics Lab, a platform for students, researchers, and industry partners to prototype ideas and test sensor systems.

Student Research Assistant

12/2015 – 03/2016

Friedrich-Alexander University, Germany

- Implemented and evaluated path-planning, mapping, and localisation methods for autonomous indoor and outdoor navigation, and helped fold results into seminar and lab teaching material.

TEACHING & SUPERVISION

Doctoral Supervision and Postgraduate Assessment

09/2024 – present

The University of Melbourne

- Co-supervising two PhD candidates alongside their primary supervisors; one submitting for examination in 2026. Research spans diffusion-based medical image generation and reconstruction, and efficient training and fine-tuning of vision state-space models. Supervision has contributed to peer-reviewed output including an ICLR 2026 paper and a reconstruction preprint.
- Marking of capstone Master's research theses in the Master of Information Technology program.

Head Teaching Associate (part-time)

02/2021 – 03/2023

Monash University, Department of Electrical and Computer Systems Engineering

- Co-designed and taught tutorials and hands-on labs on Computer Vision and on Neural Networks and Deep Learning at undergraduate and postgraduate level, in person and online, teaching 400+ students per year on topics from edge detection to CNN-based image segmentation.
- Marked student assessments and co-managed the marking across the teaching team.
- Led weekly teaching meetings, managed workload distribution, and resolved day-to-day teaching issues across a team of 7 to 12 teaching associates.
- Held one-on-one sessions supporting students through course difficulties.

Supervision and Lecturing as Research & Teaching Associate

05/2017 – 12/2019

Friedrich-Alexander University, Germany

- Co-developed and taught, as lecturer, a hands-on postgraduate Mobile Robotics lab course, consistently rated top subject (highest band) by students.
- Co-supervised 12 under- and postgraduate students to successful thesis completion, largely on thesis topics I devised spanning computer vision, probabilistic state estimation, and robotics; separately managed two student research assistants.
- Co-developed a postgraduate research seminar and co-taught Integrated Navigation Systems to under- and postgraduate students.

Student Teaching Assistant (part-time)

10/2014 – 02/2016

Friedrich-Alexander University, Germany

- Tutored introductory Control Theory for undergraduate students.

PROFESSIONAL EXPERIENCE

Software Engineer (Competition and Internship)

09/2015 – 09/2016

Audi AG, Germany, Pre-development for autonomous low-speed driving

- Developed computer vision algorithms for 3D and RGB-D camera-based environment perception and obstacle avoidance for autonomous low-speed driving on a 1:8 model car, winning the Audi Autonomous Driving Cup with team FAUtonOHM.
- Co-designed a state-based software architecture for event-driven navigation and dynamic rerouting of the model car through a complex, city-like course, with reliable failure handling and recovery.
- Ran timing analysis of the object-detection sensor chain for a real-world prototype and supported a live demo of Audi's autonomous valet parking service (internship from 04/2016).

AWARDS & RECOGNITION

- Gold Reviewer, ICML 2026, highest reviewer status, with complimentary registration.
- Outstanding Reviewer, ICLR 2024, with complimentary registration.

- Best Reviewer, ICML 2024, announced at the conference.
- Top Reviewer, NeurIPS 2023, with complimentary registration.
- Selected for the Machine Learning Summer School (MLSS) 2026, Melbourne (competitive admission).
- Melbourne Research Scholarship, full PhD scholarship with tuition offset (2021 to 2025).
- Monash Research Scholarship, full PhD scholarship with tuition offset (2020 to 2021).
- Winner, Audi Autonomous Driving Cup, first place, 10,000 EUR prize (2016).
- German Federal Scholarship (Deutschlandstipendium), 2015 and 2016.

ACADEMIC SERVICE & OUTREACH

- Conference reviewing: NeurIPS (since 2022), ICLR (since 2023), ICML (since 2023), ACCV (2022).
- Journal reviewing: IEEE TPAMI (since 2023), TMLR (since 2026).
- Earlier robotics reviewing: IROS (2019 to 2022), RA-L (2022).
- Conference volunteering: EACL (2023).
- Public outreach: co-organised an interactive robotics demonstration for non-technical audiences at the Long Night of Sciences (Germany, 2019).

PROFESSIONAL DEVELOPMENT

- Melbourne Academy, University of Melbourne, completed (2026); career development program on academic expectations in education, research, and leadership for early-career academics.
- Emerging Research Leaders Certificate, University of Melbourne, in progress (2026); a research leadership development program for early-career researchers.
- Certificate in Small-Group Teaching (CSGT), University of Melbourne, to be completed in 2026.

SELECTED PUBLICATIONS

Preprints

Zhenkai Zhang, **Markus Hiller**, Krista A. Ehinger, Tom Drummond. From Sparse X-rays to 3D CT: Training-Free Reconstruction with Diffusion Priors. [arXiv:2606.20763](https://arxiv.org/abs/2606.20763).

Cho Lei, Yanbei Jiang, **Markus Hiller**, Zhijian Zhou, Xunye Tian, Krista A. Ehinger, Nir Lipovetzky. SVoT: State-aware Visualization-of-Thought for Spatial Reasoning via Reinforcement Learning. [arXiv:2606.11770](https://arxiv.org/abs/2606.11770).

Published

Zhenkai Zhang, **Markus Hiller**, Krista A. Ehinger, Tom Drummond. Pixel-Level Residual Diffusion Transformer: Scalable 3D CT Volume Generation. International Conference on Learning Representations (ICLR), 2026.

Viktoria Schram, **Markus Hiller**, David Beck, Trevor Cohn. Active Budget Allocation for Efficient Scaling Law Estimation via Surrogate-Guided Pruning. International Conference on Machine Learning (ICML), 2026.

Viktoria Schram, **Markus Hiller**, David Beck, Trevor Cohn. Zero-Shot Performance Prediction for Probabilistic Scaling Laws. Advances in Neural Information Processing Systems 38 (NeurIPS), 2025.

Markus Hiller, Krista A. Ehinger, Tom Drummond. Perceiving Longer Sequences with Bi-Directional Cross-Attention Transformers. Advances in Neural Information Processing Systems 37 (NeurIPS), 2024.

Markus Hiller, Mehrtash Harandi, Tom Drummond. On Enforcing Better-Conditioned Meta-Learning for Rapid Few-Shot Adaptation. Advances in Neural Information Processing Systems 35 (NeurIPS), 2022.

Markus Hiller, Rongkai Ma, Mehrtash Harandi, Tom Drummond. Rethinking Generalization in Few-Shot Classification. Advances in Neural Information Processing Systems 35 (NeurIPS), 2022.

Tianyu Zhu, **Markus Hiller**, Mahsa Ehsanpour, Rongkai Ma, Tom Drummond, Ian Reid, Hamid Rezatofighi. Looking Beyond Two Frames: End-to-end Multi-object Tracking using Spatial and Temporal Transformers. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2022.

Andreas Blank, **Markus Hiller**, Siyi Zhang, et al.. 6DoF Pose-estimation Pipeline for Texture-less Industrial Components in Bin Picking Applications. European Conference on Mobile Robots (ECMR), 2019.

Markus Hiller, Chen Qiu, Florian Particke, Christian Hofmann, Jörn Thielecke. Learning Topometric Semantic Maps from Occupancy Grids. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2019.

Full publication list, including earlier robotics and object-tracking work since 2017, available on Google Scholar.

TECHNICAL SKILLS

Python, PyTorch; distributed multi-GPU and multi-node training; Slurm (HPC); Git; Docker; Kubernetes (client); C++/CUDA; WandB; Linux; AI-assisted coding tools (Claude Code, Cursor, Antigravity).

RESEARCH INTERESTS

Efficient and low-resource computer vision; geometry of optimisation; continual learning; efficient neural architecture design; retrieval-integrated generative modelling (vision and language); medical image generation and reconstruction; multimodal and vision-language models.

REFEREES

Available on request.