Lucas Beyer

Staff Research Scientist at Google Brain/DeepMind Zürich
co-leading a multimodal research effort.

Personal information

Name:	Dr. Ing. Lucas Beyer
Birth and citizenship:	Born in 1988, of French and German citizenship, grew up in Belgium.

Education

06.2013 – 04.2018	Ph.D. student at the Visual Computing Institute (VCI), RWTH Aachen University (Germany) Topic: Deep Visual Human Sensing with Application in Robotics Supervisor: Prof. Dr. Bastian Leibe Status: Defended with magna cum laude
11.2012 – 04.2013	Ph.D. student at the Aachen Institute for Advanced Study in Computational Engineering Sciences (AICES, RWTH Aachen) Subject: High-performance Density Functional Theory Supervisor: Prof. Paolo Bientinesi, Ph.D. Status: Canceled and switched topic
09.2006 – 07.2012	Diploma (eq. M.Sc.) student at the RWTH Aachen University (Germany), including a NRW-scholarship for excellent studies. Studies: Computational Engineering Sciences (CES), Grade 1.3 (A, "Very good") NRW-Scholarship: 2010 – 2011 Project thesis: Data-based modeling of protein-protein interaction, Grade 1.3 (A) Diploma thesis: Exploiting Graphics Accelerators for Computational Biology, Grade 1.0 (A)
09.1999 – 06.2006	Athénée César Franck (Belgium): High school. Grade: 2.5 (Good, B) Acquired grade: High school diploma, Grade 2.5 (B, "Good") Emphasis: Math and science

Employment history

07.2018 – present	**Staff Research Scientist:** Google Brain/DeepMind, Zürich, Switzerland Performing fundamental research in multimodal AI. Many publications, see below.
06.2017 – 09.2017	**Intern (research)**: Google, Venice, Los Angeles Disentangling representations learned by FaceNet to improve prediction tasks (go/2fn).
08.2016 – 11.2016	**Intern (AI)**: Kindred.ai, Toronto Work on robot's learning to solve tasks demonstrated by humans.
05.2016 – 08.2016	**Intern (research)**: Google, Venice, Los Angeles Neural network to determine what people in an image are looking at (go/image-gaze).
02.2012 – 11.2012	**Student assistant**: AICES, RWTH Aachen Univ. Writing a library to use a virtual reality 3D projector and sensors

10.2011 – 12.2011	**Intern (programmer)**: Mint medical GmbH, Heidelberg Implemented a graph-cut based segmentation algorithm.
03.2011 – 09.2011	**Student assistant**: LFB, RWTH Aachen Univ. Computer-assisted diagnosis for early stage pleural mesothelioma
10.2010 – 02.2011	**Tutor**: SWC, RWTH Aachen Univ. „Systematic Software Engineering“
11.2008 – 07.2009	**Tutor**: AVT & CATS, RWTH Aachen Univ. „Simulation technology 2“ and „Simulation technology for mechanical engineers“
04.2008 – 07.2008	**Tutor**: Computer Graphics, RWTH Aachen Univ. „Data-structures and algorithms“
11.2007 – 02.2008	**Tutor**: Languages & Verification, RWTH Aachen Univ. „Introduction to programming“
09.2007, 08.2008 02.2007 – 04.2007 07.2006 – 09.2006	**Programmer**: Digatron Industrie-Elektronik GmbH Software developer (C++, MFC, .Net, MSSQL) Designed and implemented new features. Designed and implemented a software from the ground up to the maturity phase.

Skills and qualifications

Languages

French:	Mother tongue
German:	Mother tongue
English:	Proficient (RWTH Aachen Univ, CEFR level C1, Grade 1.7)
Dutch and Thai:	Basic spoken/written and spoken knowledge, respectively.

Programming Languages

Proficient in C++	Around 15 years of experience in open-source, industry and academia. In-depth knowledge of the STL, OOP, design patterns, template metaprogramming. Not up-to-date beyond C++14.
Proficient in Python	Around 15 years of experience in open-source, industry and academia. Deep understanding of the LISP-y parts and, of course, duck-typing.
Good knowledge of	{Java|Coffee}Script and HTML+CSS.
Past hands-on experience:	Julia, Go, PHP, D, Clojure, Mathematica, Matlab, Prolog, ObjectiveC and unfortunately Java, .Net, *SQL

Frameworks and Tools

Tooling	Arch, Vim, Git/Hg, SVN, {C,}Make, GCC, msvc, valgrind, fish, tmux
Number crunching & “AI”	Jupyter/colab, TPUs, Jax/Theano/TensorFlow/Torch/PyTorch, NumPy/SciPy, BLAS/LAPACK, OpenMP, CUDA, OpenCL, Julia
Writing and publishing	LaTeX, matplotlib
Gamedev	OpenGL (v2/3, GLSL), SDL, SFML, OpenAL, Löve, Blender
Application development	Qt, MFC, .Net
Webdev	Docker, CherryPy, MongoDB, MySQL (MariaDB), Nginx, Apache.

Selected Publications

* indicates equal contribution

These are just a handful of selected publications, with a recency bias. For the full list, please see my Google Scholar page. As of June 2024, I have over 50k citations across over 50 papers, of which 20 have more than 100 citations.

PaliGemma: A versatile 3B VLM for transfer (tech report)
Lucas Beyer*, Andreas Steiner*, André Susano Pinto*, Alexander Kolesnikov*, Xiao Wang*, +29, Xiaohua Zhai*

No Filter: Cultural and Socioeconomic Diversity in Contrastive Vision-Language Models (under submission)
Angéline Pouget, **Lucas Beyer**, +3, Xiaohua Zhai, Ibrahim Alabdulmohsin

Getting vit in shape: Scaling laws for compute-optimal model design (NeurIPS'23)
Ibrahim Alabdulmohsin*, Xiaohua Zhai*, Alexander Kolesnikov, **Lucas Beyer***

Image captioners are scalable vision learners too (NeurIPS'23 **Oral**)
Michael Tschannen*, Manoj Kumar*, Andreas Steiner*, Xiaohua Zhai, Neil Houlsby, **Lucas Beyer***

Tuning computer vision models with task rewards (ICML'23)
André Susano Pinto*, Alexander Kolesnikov*, Yuge Shi, **Lucas Beyer**, Xiaohua Zhai

Sigmoid loss for language image pre-training (ICCV'23 **Oral**)
Xiaohua Zhai*, Basil Mustafa, Alexander Kolesnikov, **Lucas Beyer***

FlexiViT: One Model for All Patch Sizes (CVPR'23)
Lucas Beyer*, Pavel Izmailov*, Alexander Kolesnikov*, Mathilde Caron*, Simon Kornblith*, Xiaohua Zhai*, +4

UViM: A unified modeling approach for vision with learned guiding codes (NeurIPS'22)
Alexander Kolesnikov*, André Susano Pinto*, **Lucas Beyer***, Xiaohua Zhai*, Jeremiah Harmsen*, Neil Houlsby*

The Efficiency Misnomer (ICLR'22)
Mostafa Dehghani*, Anurag Arnab*, **Lucas Beyer***, Ashish Vaswani, Yi Tay*

How to train your ViT? Data, Augmentation, and Regularization in Vision Transformers
(TMLR **Outstanding paper finalist**)
Andreas Steiner*, Alexander Kolesnikov*, Xiaohua Zhai*, Ross Wightman, Jakob Uszkoreit, **Lucas Beyer***

Knowledge distillation: A good teacher is patient and consistent (CVPR'22)
Lucas Beyer*, Xiaohua Zhai*, Amélie Royer*, Larisa Markeeva*, Rohan Anil, Alexander Kolesnikov*

Scaling Vision Transformers (CVPR'22)
Xiaohua Zhai*, Alexander Kolesnikov*, Neil Houlsby, **Lucas Beyer***

An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale (ICLR'21, **Oral**)
Alexey Dosovitskiy*, **Lucas Beyer***, Alexander Kolesnikov*, Dirk Weissenborn*, Xiaohua Zhai*, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, Neil Houlsby*

MLP-Mixer: An all-MLP Architecture for Vision (NeurIPS'21)
Ilya Tolstikhin*, Neil Houlsby*, Alexander Kolesnikov*, **Lucas Beyer***, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, Alexey Dosovitskiy

Big Transfer (BiT): General Visual Representation Learning (ECCV'20 **Spotlight**)
Alexander Kolesnikov*, **Lucas Beyer***, Xiaohua Zhai*, Joan Puigcerver, Jessica Yung, Sylvain Gelly, Neil Houlsby

Are we done with ImageNet? (ArXiv)
Lucas Beyer*, Olivier J Hénaff*, Alexander Kolesnikov*, Xiaohua Zhai*, Aäron van den Oord*

S^4L: Self-supervised Semi-supervised Learning (ICCV'19 **Oral,** ICML'19 Workshop)
Xiaohua Zhai*, Avital Oliver*, Alexander Kolesnikov*, **Lucas Beyer***

Revisiting Self-supervised Visual Representation Learning (CVPR'19, ICML'19 Workshop **Oral**)
Alexander Kolesnikov*, Xiaohua Zhai*, **Lucas Beyer***

Biternion Nets: Continuous Head Pose Regression from Discrete Training Labels (GCPR'15, **Oral**)
Lucas Beyer, Alexander Hermans, Bastian Leibe

Teaching Experience

Upcoming: 07.2024 M2L Summer School **lecture** (upcoming)
08.2023 RANLP Summer School **lecture** (page)
06.2023 ACDL Summer School **lecture** (page)
07.2022 CIFAR DLRL Summer School **lecture** (page)
10.2022 M2L Summer School **lecture** (page, slides tweet, recording, feedback)
10.2021 Stanford CS25 Transformers **lecture** (page, recording)
2015-2017 RWTH Univ. Advanced Machine Learning **class TA and co-designer** of course
2015 RWTH Univ. Image Processing **seminar** (co-organizer and supervisor)
2013 RWTH Univ. 3D Computer Vision with Kinect **seminar** (co-organizer and supervisor)
2012 RWTH Univ. High-Perf. and Sci. Computing **seminar** (co-organizer and supervisor)
2012 RWTH Univ. Languages for Scientific Computing **class TA**
2010 RWTH Univ. Systematic Software Engineering **tutor**
2009 RWTH Univ. Simulation Technology for Mechanical Engineers **tutor**
2008 RWTH Univ. Simulation Technology 2 **tutor**
2008 RWTH Univ. Data-structures and Algorithms **tutor**
2007 RWTH Univ. Introduction to Programming **tutor**

Invited Talks (excludes paper orals)

Upcoming: 11.2024 ML in PL, talk (..., 7-10 Nov)

07.2024 ICML DMLR **workshop keynote + panelist** (announcement, recording)

07.2024 ELLIS **workshop talk** (announcement)

04.2024 Zürich Computer Vision **meetup talk** (announcement)

12.2023 Zürich NLP **meetup talk** (announcement)

12.2023 Co:Here 4 AI **public talk** (recording)

11.2023 Aleksa Discord **public talk** (recording)

08.2023 RANLP **keynote** (announcement)

01.2023 Co:Here 4 AI **public talk** (recording)

12.2022 NeurIPS Meta-Learning **workshop talk** (announcement)

10.2022 MLT.ai talk (announcement >400 attendees, tweet, recording)

10.2022 NittyGritty.ai (AlephAlpha but public) **public talk** (announcement, summary)

10.2022 Munich NLP **meetup talk** (recording)

09.2022 Transformers for Environmental Sciences **keynote** (announcement)

07.2022 ICML shift happens **workshop keynote** + **panelist** (announcement, feedback, thanks)

06.2022 CVPR Transformers 4 Vision **workshop panelist** (announcement)

03.2022 ZHAW **colloquium talk** (announcement)

03.2022 London ML **meetup talk** (announcement, recording)

03.2022 ByteDance invited speaker series **private talk**

10.2021 EPFL **colloquium talk** (recording)

07.2021 HuggingFace Jax/Flax **hack-week talk** (announcement, recording)

11.2017 DeepMind **private talk**

11.2017 Microsoft Research Cambridge **private talk**

08.2016 Guelph University **private talk**

Academic Community Service

Workshop organizer:

- **NeurIPS 2021:** ImageNet past present future
- **CVPR 2022-2024:** Transformers 4 Vision 1/2/3
- **Google DeepMind 2024:** internal workshop on multimodal research

CIFAR reviewer for pan-Canadian AI strategy.

Area Chair:

- NeurIPS 2024
- NeurIPS 2023 (Notable AC, top 8.1%)
- ICCV 2023

Reviewer: for NeurIPS (2021, 2019), CVPR (2021, 2020, 2019, 2018, 2017, 2016), ECCV/ICCV (2019, 2018, 2017, 2016), ICRA (2018, 2016), IROS (2016), RA-L (2019), WACV (2021), CVIU (2016), Various workshops. Multiple **outstanding reviewer awards**

Mentoring: EEML Summer School mentor, Google Research interns mentor, CIFAR DLRL faculty mixer.

Hobbies and interests

Since 1997	Playing ice hockey in Aachen (Ger), Eschweiler (Ger), Liège (Bel), and RWTH Aachen University (Ger).
03.2010 – 10.2011	Coach of the RWTH Aachen University ice-hockey team.
Since ~2000	Programming of open-source and hobby software.

Other misc. Achievements

- Held the state-of-the-art ("#1 worldwide") in most academic topics I published.
- Winner of the Google Developer Group Aachen hackathon. (As a team of 2.)
- Finished top-10% in the National Data Science Bowl machine-learning competition.
- Awarded "Bildungsfonds" scholarship for most promising students.
- AICES doctoral fellowship for extremely well-qualified students.
- Awarded Think-cell student grant to visit Meeting-C++ 2014.
- Visited Machine Learning Summer School 2014.
- Visited Deep Learning Summer School 2015. (Acceptance rate: 30%)
- Was coach of the University's ice-hockey team for two seasons.