Tim Cooijmans

cooijmans.tim@gmail.com https://cooijmanstim.github.io I am an ML&AI researcher interested in the **dynamics of learning systems**, including **meta-learning** and **self-improvement**. *Recurrent Batch Normalization* stabilized the inference of recurrent neural networks, dramatically improving training and generalization. I am currently working in the area of **multi-agent reinforcement learning**, where gradient descent fails in a way that I think can teach us something about learning in general. *Meta-Value Learning* renders differentiable games amenable to gradient descent by modeling the learning process.

Employment history

- **DeepMind** (London, UK): research internship with James Martens (2017)
 - Designed variance reduction techniques for approximate forward-mode autodiff
 - Contributed code to the TensorFlow library for HPC&ML (see Open-source contributions)
- Google Brain (Mountain View, USA): research internship with Douglas Eck & Fred Bertsch (2016)
 - Conceived Waybackprop: BPTT in O(log n) space (blog post)
 - Conceived a generative model of Bach's chorales that later powered Google's first AI Doodle
 - Contributed code to the TensorFlow library for HPC&ML (see Open-source contributions)
- Mila (Montréal, Canada): research internship with Aaron Courville & Yoshua Bengio (2015-2016)
 - Research into deep neural networks, specifically dynamic capacity networks, recurrent attention for video and batch normalization in recurrent neural networks (see *Publications*)
 - Contributed code to the Theano library for HPC&ML (see Open-source contributions)
- CERN (Geneva, Switzerland): technical internship (2014-2015)
 - Developed tools around the Monte-Carlo radsim package FLUKA to predict radioactivity in Large Hadron Collider experiment caverns
 - Trained in radiation safety, first aid & the use of fire extinguishers & oxygen self-rescue masks
- Ideaspool (Maastricht, The Netherlands): software development (2011-2012, part-time)
 - Helped develop an on-line education platform
 - Developed a smartphone app for the academic hospital
- Mapscape (Eindhoven, The Netherlands): software development (2008-2009)
 - Developed tools to manage POI data sources
 - Compram (Eindhoven, The Netherlands): software development (2006-2008)
 - Equipped time-tested ERP product from the 80s with a web interface
 - Assumed system administration and other all-round technical work
 - NoxLogic (Gorinchem, The Netherlands): software development (2005)
 - Designed and implemented a billing system in PHP+MySQL
 - Rewrote a bottleneck PHP script of the main webshop product as an Apache module in C

Education history

- PhD in Computer Science at Mila Quebec AI Institute, Université de Montréal (2016-2024)
 - Research in generative models, recurrent neural nets, credit assignment, differentiable games
 - Reviewed for ICLR, ICML and NeurIPS
 - Internships at Brain and DeepMind (see *Employment history*)
- MSc in Operations Research at Maastricht University in The Netherlands (2013-2016)

- Cum laude, GPA 8.53 (out of 10)
- Two 12-month internships at CERN and Mila (see *Employment history*)
- BSc in Knowledge Engineering at Maastricht University in The Netherlands (2009-2013)
 - GPA 8.06 (out of 10)
 - Projects completed in the context of problem-based learning:
 - design and implementation of a realistic billiards simulator and an AI to play it
 - implementation of the board game *Ticket to Ride* and several AI strategies
 - design and implementation of a traffic simulator to study traffic light control
 - application of deep neural nets and reinforcement learning to Nao robot locomotion
 - Dissertation in high-dimensional statistics that won a Thesis Award

Publications

- Advantage Alignment Algorithms: simple, cheap, effective opponent shaping
 Juan Agustin Duque, Milad Aghajohari, Tim Cooijmans, Razvan Ciuca, Tianyu Zhang, Gauthier Gidel, Aaron Courville. Submitted to ICLR 2025.
- Coordination in Generative Modeling, Automatic Differentiation and Multi-Agent Learning Tim Cooijmans. Phd dissertation, 2024.
- ReaLchords: tailoring generative models for jamming with RL fine-tuning and distillation Yusong Wu, Tim Cooijmans, Kyle Kastner, Adam Roberts, Ian Simon, Alexander Scarlatos, Chris Donahue, Cassie Tarakajian, Shayegan Omidshafiei, Aaron Courville, Pablo Samuel Castro, Natasha Jaques, Cheng-Zhi Anna Huang. ICML 2024.
- LOQA: solving social dilemmas efficiently by shaping opponent Q-values Milad Aghajohari, Juan Augustin Duque, Tim Cooijmans, Aaron Courville. ICLR 2024.
- Meta-Value Learning: solving social dilemmas with a novel general meta-learning framework Tim Cooijmans, Milad Aghajohari, Aaron Courville. Submitted to TMLR.
- Best-Response Shaping: solving social dilemmas by differentiating through the best response Milad Aghajohari, Tim Cooijmans, Juan Augustin Duque, Shunichi Akatsuka, Aaron Courville. Submitted to TMLR.
- SUNMASK: Mask-Enhanced Control in Step-Unrolled Denoising Autoencoders Kyle Kastner, Tim Cooijmans, Yusong Wu, Aaron Courville. EvoMUSART 2023.
- MIDI-DDSP: detailed **control** of **musical performance** via **hierarchical** modeling Yuson Wu, Ethan Manilow, Yi Deng, Rigel Swavely, Kyle Kastner, Tim Cooijmans, Aaron Courville, Anna Huang, Jesse Engel. ICLR 2022.
- On the Variance of UORO: an analysis of variance (reduction) for approx forward-mode autodiff Tim Cooijmans, James Martens. 2018.
- Harmonic Recomposition using **Conditional Autoregressive** Modeling Kyle Kastner, Rithesh Kumar, Tim Cooijmans, Aaron Courville. ICML 2018 workshop.
- Coconet: the ML model behind today's Bach Doodle Anna Huang, Tim Cooijmans, Monica Dinculescu, Adam Roberts, Curtis Hawthorne. Blog post, 2018.
- Memorization in Recurrent Neural Networks Tegan Maharaj, David Krueger, Tim Cooijmans. ICML 2017 PADL workshop.
- Counterpoint by Convolution: a **convolutional** model of Bach's chorales and **Gibbs sampling** strategy Anna Huang, Tim Cooijmans, Adam Roberts, Aaron Courville, Douglas Eck. ISMIR 2017.
- Waybackprop: unlocking long-term dependencies with BPTT using only logarithmic space Tim Cooijmans. Blog post, 2017.
- Recurrent Batch Normalization: the first successful use of normalization in RNN transitions Tim Cooijmans, Nicolas Ballas, César Laurent, Çağlar Gülçehre, Aaron Courville. ICLR 2017.
- Dynamic Capacity Networks: models that **dynamically allocate capacity** to salient features Amjad Almahairi, Nicolas Ballas, Tim Cooijmans, Yin Zheng, Hugo Larochelle, Aaron Courville. ICML 2016.

- Theano: a Python framework for fast computation of mathematical expressions Theano development team. 2016.
- **SESAME: radiation simulation** toolkit to modify geometry between irradiation and decay stages Tim Cooijmans, Moritz Guthoff. CERN EDMS, 2015.
- Monte-Carlo simulations of the radiation environment for the CMS experiment
 Sophie Mallows, Igor Azhgirey, Igor Bayshev, Ida Bergstrom, Tim Cooijmans, Anne Dabrowski, Lisa Glöggler, Moritz Guthoff, Igor
 Kurochkin, Helmut Vincke, S Tajeda. Elba2015 13th Pisa Meeting on Advanced Detectors.
- Technical proposal for the phase-II upgrade of the CMS detector CMS collaboration. CERN EDMS, 2015.

Open-source contributions

- Meta-Value Learning code repository
- Recurrent Batch Normalization code repository
- MCTS chess engine in C++ (2014)
- TensorFlow: add stop_gradients argument to tf.gradients
- TensorFlow: fix race condition in tf.nn.moments
- Theano: computation graph walker that works across scan boundaries
- Theano: implement various batched products
 - batched_tensordot without Scan
 - batched_dot
 - add CUDA streams path to batched_dot
 - sharded batched GEMM
 - batched GEMM in gpuarray backend
- Rclone: enable Google Drive from service accounts
- Ratpoison: make keymaps switchable at runtime