

Mengye Ren – Curriculum Vitæ

CONTACT INFORMATION	60 5th Ave, Rm 508 New York, NY 10011-8868, USA	<i>Tel:</i> +1 (212) 992-7547 <i>Email:</i> mengye@nyu.edu <i>Website:</i> https://mengyeren.com
RESEARCH INTERESTS	General areas: machine learning, computer vision, representation learning, continual learning, meta-learning, few-shot learning, artificial intelligence. My key research question is: how do we enable human-like, agent-based machine intelligence to continually learn, adapt, and reason in naturalistic environments? I am interested in the emergence of intelligence by learning from a point-of-view experience.	
EDUCATION	University of Toronto <i>Department of Computer Science</i> Ph.D., Supervisors: Richard Zemel and Raquel Urtasun Thesis: “Open World Machine Learning with Limited Labeled Data” Committee: Roger Grosse, Geoffrey Hinton, Yee Whye Teh <i>Department of Computer Science</i> M.Sc., Supervisor: Richard Zemel <i>Department of Engineering Science, ECE Option</i> B.A.Sc., CGPA 3.90/4.00, with high honours	2017/01 – 2021/10 2015/09 – 2017/01 2011/09 – 2015/06
PROFESSIONAL EXPERIENCE	New York University , New York, NY, USA <i>Assistant Professor of Computer Science and Data Science</i> Google Brain , Toronto, ON, Canada <i>Visiting Faculty Researcher</i> Waabi Innovation , Toronto, ON, Canada <i>Senior Researcher II</i> Uber ATG , Toronto, ON, Canada <i>Senior Research Scientist I</i> <i>Research Scientist II</i> Twitter , Cambridge, MA, USA <i>Research Intern</i> Google , Mountain View, CA, USA <i>Research Intern in Google Brain and Image Understanding</i> <i>SWE Intern in StreetView</i> Microsoft , Redmond, WA, USA <i>SDET Intern in Visual Studio</i> <i>SDET Intern in Visual Studio</i>	2022/09 – Present 2022/01 – 2022/09 2021/03 – 2021/12 2018/09 – 2021/02 2017/05 – 2018/09 2016/06 – 2016/08 2015/06 – 2015/08 2014/06 – 2014/08 2013/06 – 2013/08 2012/06 – 2012/08
TEACHING	New York University <ul style="list-style-type: none">DS-GA 3001: Embodied Learning and VisionDS-GA 3001: Embodied Learning and VisionCSCI-GA 2565: Machine LearningDS-GA 1008: Deep LearningCSCI-GA 2565: Machine Learning	2026 Spring 2025 Spring 2024 Fall 2024 Spring 2023 Fall

- DS-GA 1003: Machine Learning 2023 Spring

Vector Institute

- Deep Learning II 2020 Fall

University of Toronto

- CSC 411: Machine Learning and Data Mining 2019 Winter

PEER-REVIEWED CONFERENCE PUBLICATIONS

(* = equal contribution)

2026

- C49 Ying Wang, Oumayma Bounou, Gaoyue Zhou, Randall Balestriero, Tim G. J. Rudner, Yann LeCun, and **Mengye Ren**. Temporal straightening for latent planning. *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, Seoul, South Korea, 2026.
- C48 Jack Lu, Ryan Teehan, Zhenbang Yang, and **Mengye Ren**. Context tuning for in-context optimization. *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, Seoul, South Korea, 2026.
- C47 Zayne Sprague, Jack Lu, Manya Wadhwa, Sedrick Keh, **Mengye Ren**, and Greg Durrett. Skill-Factory: Self-distillation for learning cognitive behaviors. *Proceedings of the 14th International Conference on Learning Representations (ICLR)*, Rio de Janeiro, Brazil, 2026.
- C46 Frank (Zequan) Wu and **Mengye Ren**. Local reinforcement learning with action-conditioned root mean squared Q-functions. *Proceedings of the 14th International Conference on Learning Representations (ICLR)*, Rio de Janeiro, Brazil, 2026.
- C45 Christopher Hoang and **Mengye Ren**. Midway Network: Learning representations for recognition and motion from latent dynamics. *Proceedings of the 14th International Conference on Learning Representations (ICLR)*, Rio de Janeiro, Brazil, 2026.

2025

- C44 Yasaman Mahdaviyeh, James Lucas, **Mengye Ren**, Andreas S. Tolias, Richard Zemel, and Toniann Pitassi. Replay can provably increase forgetting. *Proceedings of the 4th Conference on Lifelong Learning Agents (CoLLAs)*, Philadelphia, Pennsylvania, USA, 2025.
- C43 Yanlai Yang and **Mengye Ren**. Memory Storyboard: Leveraging temporal segmentation for streaming self-supervised learning from egocentric videos. *Proceedings of the 4th Conference on Lifelong Learning Agents (CoLLAs)*, Philadelphia, Pennsylvania, USA, 2025. (oral)
- C42 Raghav Singhal*, Zachary Horvitz*, Ryan Teehan*, **Mengye Ren**, Zhou Yu, Kathleen McKeown, and Rajesh Ranganath. A general framework for inference-time scaling and steering of diffusion models. *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, Vancouver, British Columbia, Canada, 2025.
- C41 Hui Dai, Ryan Teehan, and **Mengye Ren**. Are LLMs prescient? A continuous evaluation using daily news as the oracle. *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, Vancouver, British Columbia, Canada, 2025.
- C40 Alex N. Wang*, Christopher Hoang*, Yuwen Xiong, Yann LeCun, and **Mengye Ren**. PooDL: Pooled and dense self-supervised learning from naturalistic videos. *Proceedings of the 13th International Conference on Learning Representations (ICLR)*, Singapore, 2025.

2024

- C39 Jack Lu, Ryan Teehan, and **Mengye Ren**. ProCreate, don't reproduce! Propulsive energy diffusion for creative generation. *Computer Vision – the 18th European Conference (ECCV)*, Milan, Italy, 2024.

- C38 Yipeng Zhang, Laurent Charlin, Richard Zemel, and **Mengye Ren**. Integrating present and past in unsupervised continual learning. *Proceedings of the 3rd Conference on Lifelong Learning Agents (CoLLAs)*, Pisa, Italy, 2024. **(oral)**
- C37 Ryan Teehan, Brenden M. Lake, and **Mengye Ren**. CoLLEGe: Concept embedding generation for large language models. *Proceedings of the 1st Conference on Language Modeling (COLM)*, Philadelphia, Pennsylvania, USA, 2024.
- C36 Yanlai Yang, Matt Jones, Michael C. Mozer, and **Mengye Ren**. Reawakening knowledge: Anticipatory recovery from catastrophic interference via structured training. *Advances in Neural Information Processing Systems 37 (NeurIPS)*, Vancouver, British Columbia, Canada, 2024.
- C35 Emin Orhan, Wentao Wang, Alex N. Wang, **Mengye Ren**, and Brenden M. Lake. Self-supervised learning of video representations from a child’s perspective. *Proceedings of the 44th Annual Meeting of the Cognitive Science Society (CogSci)*, Rotterdam, the Netherlands, 2024.
- C34 Jiachen Zhao, Zhun Deng, David Madras, James Zou, and **Mengye Ren**. Learning and forgetting unsafe examples in large language models. *Proceedings of the 41st International Conference on Machine Learning (ICML)*, Vienna, Austria, 2024.

2023

- C33 Lunjun Zhang, Anqi Joyce Yang, Yuwen Xiong, Sergio Casas, Bin Yang, **Mengye Ren**, and Raquel Urtasun. Towards unsupervised object detection from LiDAR point clouds. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, British Columbia, Canada, 2023.
- C32 **Mengye Ren**, Simon Kornblith, Renjie Liao, and Geoffrey Hinton. Scaling forward gradient with local losses. *Proceedings of the 11th International Conference on Learning Representations (ICLR)*, Kigali, Rwanda, 2023.
- C31 Matt Jones, Tyler R. Scott, **Mengye Ren**, Gamaleldin F. Elsayed, Katherine Hermann, David Mayo, and Michael C. Mozer. Learning in temporally structured environments. *Proceedings of the 11th International Conference on Learning Representations (ICLR)*, Kigali, Rwanda, 2023.
- C30 David Mayo, Tyler R Scott, **Mengye Ren**, Gamaledin Elsayed, Katherine Hermann, Matt Jones, and Michael Mozer. Multitask learning via interleaving: A neural network investigation. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society (CogSci)*, Sydney, Australia, 2023.

2022

- C29 Chris Zhang*, Runsheng Guo*, Wenyuan Zeng*, Yuwen Xiong, Binbin Dai, Rui Hu, **Mengye Ren**, and Raquel Urtasun. Rethinking closed-loop training for autonomous driving. *Computer Vision – the 17th European Conference (ECCV)*, Tel-Aviv, Israel, 2022.

2021

- C28 Sean Segal*, Nishanth Kumar*, Sergio Casas, Wenyuan Zeng, **Mengye Ren**, Jingkan Wang, and Raquel Urtasun. Just label what you need: Fine-grained active selection for perception and prediction through partially labeled scenes. *Conference on Robot Learning (CoRL)*, London, United Kingdom, 2021.
- C27 James Tu, Huichen Li, Xinchun Yan, **Mengye Ren**, Yun Chen, Ming Liang, Eilyan Bitar, Ersin Yumer, and Raquel Urtasun. Exploring adversarial robustness of multi-sensor perception systems in self driving. *Conference on Robot Learning (CoRL)*, London, United Kingdom, 2021.
- C26 James Tu*, Tsun-Hsuan Wang*, Jingkan Wang, Sivabalan Manivasagam, **Mengye Ren**, and Raquel Urtasun. Adversarial attacks on multi-agent communication. *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- C25 Yuwen Xiong, **Mengye Ren**, Wenyuan Zeng, and Raquel Urtasun. Self-supervised representation learning from flow equivariance. *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

- C24 Jingkang Wang, Ava Pun, James Tu, Abbas Sadat, Sergio Casas, Sivabalan Manivasagam, **Mengye Ren**, and Raquel Urtasun. AdvSim: Generating safety-critical scenarios for self-driving vehicles. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- C23 Shuhan Tan*, Kelvin Wong*, Shenlong Wang, Sivabalan Manivasagam, **Mengye Ren**, and Raquel Urtasun. SceneGen: Learning to simulate realistic traffic scenes. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- C22 Bob Wei*, **Mengye Ren***, Wenyuan Zeng, Ming Liang, Bin Yang, and Raquel Urtasun. Perceive, attend and drive: Learning spatial attention for safe self-driving. *IEEE International Conference on Robotics and Automation (ICRA)*, Xi'an, China, 2021. (oral)
- C21 James Lucas, **Mengye Ren**, Irene Kameni, Toniann Pitassi, and Richard S. Zemel. Theoretical bounds on estimation error for meta learning. *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2021.
- C20 Alexander Wang*, **Mengye Ren***, and Richard S. Zemel. SketchEmbedNet: Learning novel concepts by imitating drawings. *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021.
- C19 **Mengye Ren**, Michael L. Iuzzolino, Michael C. Mozer, and Richard S. Zemel. Wandering within a world: Online contextualized few-shot learning. *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2021.

2020

- C18 Nicholas Vadivelu, **Mengye Ren**, James Tu, Jingkang Wang, and Raquel Urtasun. Learning to communicate and correct pose errors. *Conference on Robot Learning (CoRL)*, Cambridge, Massachusetts, USA, 2020.
- C17 Abbas Sadat*, Sergio Casas Romero*, **Mengye Ren**, Xinyu Wu, Pranaab Dhawan, and Raquel Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. *Computer Vision – the 16th European Conference (ECCV)*, Glasgow, United Kingdom, 2020.
- C16 Lingyun (Luke) Li, Bin Yang, Ming Liang, Wenyuan Zeng, **Mengye Ren**, Sean Segal, and Raquel Urtasun. End-to-end contextual perception and prediction with interaction transformer. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, Nevada, USA, 2020. (oral)
- C15 Yuwen Xiong, **Mengye Ren**, and Raquel Urtasun. LoCo: Local contrastive representation learning. *Advances in Neural Information Processing Systems 33 (NeurIPS)*, Vancouver, British Columbia, Canada, 2020.
- C14 Quinlan Sykora*, **Mengye Ren***, and Raquel Urtasun. Multi-agent routing value iteration network. *Proceedings of the 37th International Conference on Machine Learning (ICML)*, Vienna, Austria, 2020.
- C13 James Tu, **Mengye Ren**, Sivabalan Manivasagam, Ming Liang, Bin Yang, Richard Du, Frank Cheng, and Raquel Urtasun. Physically realizable adversarial examples for LiDAR object detection. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, Washington, USA, 2020.

2019

- C12 Kelvin Wong, Shenlong Wang, **Mengye Ren**, Ming Liang, and Raquel Urtasun. Identifying unknown instances for autonomous driving. *Conference on Robot Learning (CoRL)*, Osaka, Japan, 2019. (spotlight)
- C11 Abbas Sadat*, **Mengye Ren***, Andrei Pokrovsky, Yen-Chen Lin, Ersin Yumer, and Raquel Urtasun. Jointly learnable behavior and trajectory planner for self-driving vehicles. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Macau, China, 2019. (oral)

C10 **Mengye Ren**, Renjie Liao, Ethan Fetaya, and Richard S. Zemel. Incremental few-shot learning with attention attractor networks. *Advances in Neural Information Processing Systems 32 (NeurIPS)*, Vancouver, British Columbia, Canada, 2019.

C9 Chris Zhang, **Mengye Ren**, and Raquel Urtasun. Graph hypernetworks for neural architecture search. *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, New Orleans, Louisiana, USA, 2019.

2018

C8 **Mengye Ren**, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. *Proceedings of the 35th International Conference on Machine Learning (ICML)*, Stockholm, Sweden, 2018.

C7 Yuhuai Wu*, **Mengye Ren***, Renjie Liao, and Roger B. Grosse. Understanding short-horizon bias in meta optimization. *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, Vancouver, British Columbia, Canada, 2018.

C6 **Mengye Ren**, Eleni Triantafillou*, Sachin Ravi*, Jake Snell, Kevin Swersky, Joshua B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. Meta-learning for semi-supervised few-shot classification. *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, Vancouver, British Columbia, Canada, 2018.

C5 **Mengye Ren***, Andrei Pokrovsky*, Bin Yang*, and Raquel Urtasun. SBNet: Sparse blocks network for fast inference. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, Utah, USA, 2018. (**spotlight**)

2017

C4 Aidan N. Gomez*, **Mengye Ren***, Raquel Urtasun, and Roger B. Grosse. The reversible residual network: Backpropagation without storing activations. *Advances in Neural Information Processing Systems 30 (NIPS)*, Long Beach, California, USA, 2017.

C3 **Mengye Ren***, Renjie Liao*, Raquel Urtasun, Fabian H. Sinz, and Richard S. Zemel. Normalizing the normalizers: Comparing and extending network normalization schemes. *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, Toulon, France, 2017.

C2 **Mengye Ren** and Richard S. Zemel. End-to-end instance segmentation with recurrent attention. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, Hawaii, USA, 2017. (**spotlight**)

2015

C1 **Mengye Ren**, Ryan Kiros, and Richard S. Zemel. Exploring models and data for image question answering. *Advances in Neural Information Processing Systems 28 (NIPS)*, Montréal, Québec, Canada, 2015.

PEER-REVIEWED 2026

WORKSHOP PAPERS W24 Xiaoyang Jiang, Yanlai Yang, Kenneth A. Norman, Brenden Lake, and **Mengye Ren**. Baby-CLIP: Continual language grounding through the experience of a single child. *CVPR 2026 Workshop on Computer Vision for Children (CV4CHL)*, Denver, Colorado, USA, 2026.

W23 Yanlai Yang, Zhuokai Zhao, Satya Narayan Shukla, Aashu Singh, Shlok Kumar Mishra, Lizhu Zhang, and **Mengye Ren**. StreamMem: Query-agnostic KV cache memory for streaming video understanding. *CVPR 2026 Workshop on Video Large Language Models (VidLLMs)*, Denver, Colorado, USA, 2026.

W22 Jack Lu*, Ryan Teehan*, Jinran Jin, and **Mengye Ren**. When does verification pay off? A closer look at LLMs as solution verifiers. *ICLR 2026 Workshop on AI with Recursive Self-Improvement*, Rio de Janeiro, Brazil, 2026.

W21 Azwar Abdulsalam, Christopher Hoang, and **Mengye Ren**. LaMo: A latent motion world model for long-horizon prediction. *ICLR 2026 the 2nd Workshop on World Models: Understanding, Modelling and Scaling*, Rio de Janeiro, Brazil, 2026.

W20 Ying Wang, Oumayma Bounou, Gaoyue Zhou, Randall Balestriero, Tim G. J. Rudner, Yann LeCun, and **Mengye Ren**. Temporal straightening for latent planning. *World Modeling Workshop*, Montréal, Québec, Canada, 2026. (oral)

2025

W19 Ellen Su*, Solim LeGris*, Todd M. Gureckis, and **Mengye Ren**. Opinion: Learning intuitive physics requires more than visual data. *NeurIPS 2025 Workshop on Embodied World Models for Decision Making*, San Diego, California, USA, 2025.

W18 Jack Lu, Ryan Teehan, Zhenbang Yang, and **Mengye Ren**. Context tuning for in-context optimization. *Second Workshop on Test-Time Adaptation: Putting Updates to the Test! at ICML 2025*, Vancouver, British Columbia, Canada, 2025.

W17 Junyeob Baek, Hosung Lee, Christopher Hoang, **Mengye Ren**, and Sungjin Ahn. Discrete JEPA: Learning discrete token representations without reconstruction. *WoTok: Workshop on Tokenization at ICML 2025*, Vancouver, British Columbia, Canada, 2025.

2024

W16 Hui Dai, Ryan Teehan, and **Mengye Ren**. Are LLMs prescient? A continuous evaluation using daily news as the oracle. *NeurIPS Adaptive Foundation Models Workshop (AFM)*, 2024.

W15 Yipeng Zhang, Laurent Charlin, Richard Zemel, and **Mengye Ren**. Integrating present and past in unsupervised continual learning. *5th CVPR Workshop on Continual Learning (CLVision)*, Seattle, Washington, USA, 2024.

2023

W14 Ying Wang, Dongdong Sun, Rui Chen, Yanlai Yang, and **Mengye Ren**. Egocentric video comprehension via large language model inner speech. *3rd International Ego4D Workshop at CVPR*, Vancouver, British Columbia, Canada, 2023.

2022

W13 Andrew J. Nam*, **Mengye Ren***, Chelsea Finn, and James L. McClelland. Learning to reason with relational abstractions. *2nd Workshop on MATH-AI at NeurIPS*, New Orleans, Louisiana, USA, 2022.

W12 Matt Jones, Tyler R. Scott, Gamaleldin F. Elsayed, **Mengye Ren**, Katherine Hermann, David Mayo, and Michael C. Mozer. Neural network online training with sensitivity to multiscale temporal structure. *Memory in Artificial and Real Intelligence Workshop at NeurIPS*, New Orleans, Louisiana, USA, 2022.

2020

W11 Jingkang Wang*, **Mengye Ren***, Ilija Bogunovic, Yuwen Xiong, and Raquel Urtasun. Cost-efficient online hyperparameter optimization. *ICML RealML Workshop*, Vienna, Austria, 2020.

W10 **Mengye Ren**, Michael L. Iuzzolino, Michael C. Mozer, and Richard S. Zemel. Wandering within a world: Online contextualized few-shot learning. *ICML Continual Learning Workshop & Lifelong Learning Workshop & Workshop on Learning in Artificial Open Worlds*, Vienna, Austria, 2020. (oral)

W9 **Mengye Ren***, Eleni Triantafillou*, Kuan-Chieh Wang*, James Lucas*, Jake Snell, Xaq Pitkow, Andreas S. Tolias, and Richard S. Zemel. Flexible few-shot learning of contextual similarity. *NeurIPS Meta-Learning Workshop*, Vancouver, British Columbia, Canada, 2020.

W8 Laleh Seyyed-Kalantari, Karsten Roth, **Mengye Ren**, Parsa Torabian, Joseph P. Cohen, and Marzyeh Ghassemi. Multi-label incremental few-shot learning for medical image pathology

classifiers. *Medical Imaging Meets NeurIPS Workshop*, Vancouver, British Columbia, Canada, 2020.

2019

- W7 Yuwen Xiong*, **Mengye Ren***, Renjie Liao, Kelvin Wong, and Raquel Urtasun. Deformable filter convolution for point cloud reasoning. *arXiv preprint arXiv:1907.13079*, . In *NeurIPS Workshop on Sets & Partitions*, Vancouver, British Columbia, Canada, 2019.
- W6 James Lucas, **Mengye Ren**, and Richard S. Zemel. Information-theoretic limitations on novel task generalization. *NeurIPS Workshop on Machine Learning with Guarantees*, Vancouver British Columbia, Canada, 2019. **(oral)**

2018

- W5 **Mengye Ren**, Renjie Liao, Ethan Fetaya, and Richard S. Zemel. Incremental few-shot learning with attention attractor networks. *NeurIPS Meta-Learning Workshop*, Montréal, Québec, Canada, 2018.
- W4 Chris Zhang, **Mengye Ren**, and Raquel Urtasun. Graph hypernetworks for neural architecture search. *NeurIPS Meta-Learning Workshop*, Montréal, Québec, Canada, 2018.

2017

- W3 Yuhuai Wu*, **Mengye Ren***, Renjie Liao, and Roger B. Grosse. Understanding short-horizon bias in meta optimization. *NIPS Meta-Learning Workshop*, Long Beach, California, USA, 2017.
- W2 **Mengye Ren**, Eleni Triantafillou*, Sachin Ravi*, Jake Snell, Kevin Swersky, Joshua B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. Meta-learning for semi-supervised few-shot classification. *NIPS Meta-Learning Workshop & Learning with Limited Data Workshop*, Long Beach, California, USA, 2017.

2015

- W1 **Mengye Ren**, Ryan Kiros, and Richard Zemel. Exploring models and data for image question answering. *ICML Deep Learning Workshop*, Lille, France, 2015. **(oral)**

PREPRINTS & TECH REPORTS

- R13 **Mengye Ren**. The self requires learning. *Preprint*, 2026.
- R12 **Mengye Ren**. Seeking the unfamiliar but memorable: Conceptual creativity as meta-learning. *arXiv preprint arXiv:2605.16477*, 2026.
- R11 Kangsan Kim, Minki Kang, Taeil Kim, Yanlai Yang, **Mengye Ren**, and Sung Ju Hwang. Memory Transfer Learning: How memories are transferred across domains in coding agents. *arXiv preprint arXiv:2604.14004*, 2026.
- R10 Kangsan Kim, Yanlai Yang, Suji Kim, Woongyeong Yeo, Youngwan Lee, **Mengye Ren**, and Sung Ju Hwang. MA-EgoQA: Question answering over egocentric videos from multiple embodied agents. *arXiv preprint arXiv:2603.09827*, 2026.
- R9 Jack Lu*, Ryan Teehan*, Jinran Jin, and **Mengye Ren**. When does verification pay off? A closer look at LLMs as solution verifiers. *arXiv preprint arXiv:2512.02304*, 2025.
- R8 Ying Wang, **Mengye Ren**, and Andrew Gordon Wilson. In-context clustering with large language models. *arXiv preprint arXiv:2510.08466*, 2025.
- R7 Ying Wang, Yanlai Yang, and **Mengye Ren**. LifelongMemory: Leveraging LLMs for answering queries in long-form egocentric videos. *arXiv preprint arXiv:2312.05269*, 2023.
- R6 Yixuan Luo, **Mengye Ren**, and Sai Qian Zhang. BIM: block-wise self-supervised learning with masked image modeling. *arXiv preprint arXiv:2311.17218*, 2023.
- R5 Renjie Liao, Simon Kornblith, **Mengye Ren**, David J. Fleet, and Geoffrey Hinton. Gaussian-Bernoulli RBMs without tears. *arXiv preprint arXiv:2210.10318*, 2022.

- R4 Andrew J. Nam*, **Mengye Ren***, Chelsea Finn, and James L. McClelland. Learning to reason with relational abstractions. *arXiv preprint arXiv:2210.02615*, 2022.
- R3 **Mengye Ren**, Tyler R. Scott, Michael L. Iuzzolino, Michael C. Mozer, and Richard Zemel. Online unsupervised learning of visual representations and categories. *arXiv preprint arXiv:2109.05675*, 2021.
- R2 **Mengye Ren***, Eleni Triantafillou*, Kuan-Chieh Wang*, James Lucas*, Jake Snell, Xaq Pitkow, Andreas S. Tolias, and Richard S. Zemel. Flexible few-shot learning of contextual similarity. *arXiv preprint arXiv:2012.05895*, 2020.
- R1 Yuwen Xiong*, **Mengye Ren***, and Raquel Urtasun. Learning to remember from a multi-task teacher. *arXiv preprint arXiv:1910.04650*, 2019.

PATENTS

- P11 Raquel Urtasun, Bob Qingyuan Wei, **Mengye Ren**, Wenyuan Zeng, Ming Liang, and Bin Yang. Systems and methods for using attention masks to improve motion planning. US 12,282,328 B2, *U.S. Patent*, 2025.
- P10 Xuanyuan Tu, Raquel Urtasun, Tsun-Hsuan Wang, Sivabalan Manivasagam, Jingkan Wang, and **Mengye Ren**. Systems and methods for training machine-learned models with deviating intermediate representations. US 12,223,734 B2, *U.S. Patent*, 2025.
- P9 Jingkan Wang, Ava Alison Pun, Xuanyuan Tu, **Mengye Ren**, Abbas Sadat, Sergio Casas, Sivabalan Manivasagam, and Raquel Urtasun. Generating autonomous vehicle testing data through perturbations and adversarial loss functions. US 12,214,801 B2, *U.S. Patent*, 2025.
- P8 Nicholas Baskar Vadivelu, **Mengye Ren**, Xuanyuan Tu, Raquel Urtasun, and Jingkan Wang. Systems and methods for mitigating vehicle pose error across an aggregated feature map. US 12,127,085 B2, *U.S. Patent*, 2024.
- P7 Raquel Urtasun, Abbas Sadat, **Mengye Ren**, Andrei Pokrovsky, Yen-Chen Lin, and Ersin Yumer. Jointly learnable behavior and trajectory planning for autonomous vehicles. US 11,755,014 B2, *U.S. Patent*, 2023.
- P6 LingYun Li, Bin Yang, Wenyuan Zeng, Ming Liang, **Mengye Ren**, Sean Segal, and Raquel Urtasun. Systems and methods for generating motion forecast data for a plurality of actors with respect to an autonomous vehicle. US 11,780,472 B2, *U.S. Patent*, 2023.
- P5 LingYun Li, Bin Yang, Ming Liang, Wenyuan Zeng, **Mengye Ren**, Sean Segal, and Raquel Urtasun. Systems and methods for generating motion forecast data for a plurality of actors with respect to an autonomous vehicle. US 11,691,650 B2, *U.S. Patent*, 2023.
- P4 Xuanyuan Tu, Sivabalan Manivasagam, **Mengye Ren**, Ming Liang, Bin Yang, and Raquel Urtasun. Systems and methods for training object detection models using adversarial examples. US 11,686,848 B2, *U.S. Patent*, 2023.
- P3 Shuhan Tan, Kelvin Ka Wing Wong, Shenlong Wang, Siva Manivasagam, **Mengye Ren**, and Raquel Urtasun. Systems and methods for simulating traffic scenes. US 11,580,851 B2, *U.S. Patent*, 2023.
- P2 Raquel Urtasun, Kelvin Ka Wing Wong, Shenlong Wang, **Mengye Ren**, and Ming Liang. Systems and methods for identifying unknown instances. US 11,475,675 B2, *U.S. Patent*, 2022.
- P1 Raquel Urtasun, **Mengye Ren**, Andrei Pokrovsky, and Bin Yang. Sparse convolutional neural networks. US 11,061,402 B2, *U.S. Patent*, 2021.

AWARDS & HONORS	• Google Research Award	2026
	• NSERC Postdoctoral Fellowship, \$90,000 CAD (declined)	2021 – 2023
	• Facebook Fellowship Finalist (91 out of 1876 PhD applicants worldwide)	2020
	• NSERC Alexander Graham Bell Scholarship, \$105,000 CAD	2018 – 2021
	• NVIDIA Research Pioneer Award	2018
	• NVIDIA Research Pioneer Award	2017
	• NIPS 2017 Travel Award \$800 USD	2017
	• ICLR 2017 Travel Award \$1,250 USD	2017
	• MLSS 2015 Kyoto Travel Support ¥140,000 JPY	2015
	• U of T Quantathon 2nd Place \$5,000 CAD	2015
	• U of T Undergraduate Mathematics Competition, Honourable Mention	2015
	• Wallberg Undergraduate Scholarship \$1,500 CAD	2014
	• International 5th place in Windward AI Challenge, 1st in U of T	2014
	• Dean’s List for all semesters in undergraduate studies	2011 – 2015
	• Entrance Scholarship from the University of Toronto \$5,000 CAD	2011
• Meritorious Award in Mathematical Contest of Modeling (MCM)	2011	
RESEARCH SPONSORS	• NSF Award 2545541 <i>CompCog: Physical Reasoning in Humans and Machines</i> (Co-PI; PI: Todd M. Gureckis)	2026 – 2029
	• Toyota Research Institute: <i>Right-Time, Right-Sized Intelligence (R2I)</i> (Co-PI; PI: Henny Admoni)	2026 – 2027
	• Google Research Award	2026 – 2027
	• Visko AI Research Gift	2026 – 2027
	• NYU-KAIST Research Collaboration	2024 – 2026
	• Korea IITP Sponsored Global Frontier AI Lab: Fundamental Research in AI	2024 – 2029
	• Google Cloud Research Credit, \$5,000	2024 – 2025
	• Center for AI Safety Compute Credit	2024
	• Microsoft Accelerating Foundation Models Research Cloud Credit, \$80,000	2023 – 2025
	• Google Cloud Research Credit, \$5,000	2022 – 2023
	PROFESSIONAL SERVICE	<i>Conference Organizer:</i>
• Conference on Lifelong Learning Agents (CoLLAs) (Associate Program Chair)		2025
• Neural Information Processing Systems (NeurIPS) (Communications Chair)		2025
<i>Area Chair:</i>		
• Conference on Neural Information Processing Systems (NeurIPS)		2026
• Conference on Lifelong Learning Agents (CoLLAs)		2026
• International Conference on Machine Learning (ICML)		2026
• International Conference on Learning Representations (ICLR)		2026
• Conference on Language Modeling (COLM)		2025 - 2026
• IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)		2025
• International Conference on Automated Machine Learning (AutoML)		2022 – 2023
• NeurIPS Workshop on Meta-Learning (MetaLearn)		2020 – 2022
<i>Workshop Organizer:</i>		
• NeurIPS 2024 Workshop on Adaptive Foundation Models (Designated Contact)		2024
• NeurIPS 2024 Workshop on Compositional Learning		2024
• ICML 2023 Localized Learning Workshop (Designated Contact)		2023
• ICLR 2023 Workshop on Scene Representations for Autonomous Driving		2023
<i>Journal Reviewer:</i>		
• Journal of Machine Learning Research (JMLR)		
• International Journal of Computer Vision (IJCV)		
• IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)		

- IEEE Transactions on Image Processing (TIP)
- IEEE Transactions on Neural Networks and Learning Systems (TNNLS)
- IEEE Transactions on Computational Imaging (TCI)
- Science Advances
- Neural Networks
- Machine Learning
- Optimization Methods and Software
- Nature Computational Science
- Science Advances

Conference Reviewer:

- Conference on Neural Information Processing Systems (NeurIPS) 2016 – 2025
- International Conference on Machine Learning (ICML) 2017 – 2025
- International Conference on Learning Representations (ICLR) 2018 – 2025
- Conference on Computer Vision and Pattern Recognition (CVPR) 2018 – 2024
- International Conference on Computer Vision (ICCV) 2019 – 2023
- European Conference on Computer Vision (ECCV) 2020 – 2026
- Conference on Lifelong Learning Agents (CoLLAs) 2022 – 2024
- International Conference on Machine Learning (ICML) Workshops 2024
- International Conference on Robotics and Automation (ICRA) 2021
- International Conference on Intelligent Robots and Systems (IROS) 2020
- Robotics: Science and Systems (RSS) 2026
- Association for the Advancement of Artificial Intelligence Conference (AAAI) 2018
- Uncertainty in Artificial Intelligence (UAI) 2018

Scientific Reviewer:

- NSF Postdoctoral Fellowship Program 2024
- ETH Zurich/Swiss National Supercomputing Center 2024

Seminar (Co-)Organizer:

- Self-Supervised Learning Weekly Seminars 2021 – 2022
- Meta-Learning Weekly Seminars 2019 – 2021
- Uber ATG R&D Weekly Paper Reading Seminars 2018 – 2019

OPEN SOURCE
SOFTWARES

- Forward-mode automatic differentiation for TensorFlow.
GitHub: <https://github.com/renmengye/tensorflow-forward-ad>
- DeepDashboard: Real-time web-based training visualizer.
GitHub: <https://github.com/renmengye/deep-dashboard>

INVITED TALKS

2026

- T59 Does your LLM agent have a self? ICML 2026 MemAgents Workshop. Rio de Janeiro, Brazil. Apr 27, 2026.
- T58 The always-learning machine. Columbia University. New York, New York, USA. Apr 9, 2026.
- T57 Lifelong concept learning. New York University CDS Seminar. New York, New York, USA. Feb 13, 2026.

2025

- T56 Lifelong concept learning. Korea Advanced Institute of Science and Technology (KAIST). Seoul, South Korea. Oct 31, 2025.
- T55 Lifelong concept learning. Seoul National University. Seoul, South Korea. Oct 30, 2025.
- T54 Lifelong concept learning. Global AI Frontiers Symposium. Seoul, South Korea. Oct 27, 2025.

T53 Lifelong concept learning. University of Calgary. Department of Electrical and Software Engineering. Calgary, Alberta, Canada. Jun 16, 2025.

T52 Lifelong concept learning. IITP AI Education Program. New York, New York, USA. Jun 2, 2025.

T51 Lifelong concept learning for machines. New York University. Department of Psychology ConCats Seminar. New York, New York, USA. Apr 4, 2025.

2024

T50 Lifelong and human-like learning in foundation models. Seoul National University. Guest Lecture at SNU AI Seminar. Seoul, South Korea. Dec 5, 2024.

T49 Lifelong and human-like learning in foundation models. Columbia University. NSF AI Institute for Artificial and Natural Intelligence. New York, New York, USA. Sept 13, 2024.

T48 Computer Vision and Deep Learning: A Primer. NYU AI Summer School. New York University. New York, New York, USA. Jun 4, 2024.

T47 Lifelong and human-like learning in foundation models. Machine Learning Seminar. Flatiron Institute. New York, New York, USA. Apr 30, 2024.

T46 Lifelong and human-like learning in foundation models. Smart Minds meet Smart Machines: AI for Science and Public Good. German Consulate General in New York. New York, New York, USA. Apr 8, 2024.

2023

T45 Lifelong learning in structured environments. American Statistical Association, Statistical Learning and Data Science Webinar. Virtual. Oct 2023.

T44 Scaling forward gradient with local losses. Baylor College of Medicine, Journal Club Invited Talk. Houston, Texas, USA. Jun 2023.

T43 Biologically plausible learning using local activity perturbation. University of British Columbia. Invited Talk. Vancouver, British Columbia, Canada. Jun 2023.

2022

T42 Meta-learning within a lifetime. NeurIPS 2022 MetaLearn Workshop, Invited Talk. New Orleans, Louisiana, USA. Dec 2022.

T41 Biologically plausible learning using local activity perturbation. NYU CDS Lunch Seminar. New York, New York, USA. Oct 2022.

T40 Visual learning in the open world. 19th Conference on Vision and Robotics (CRV), Invited Symposium. Toronto, Ontario, Canada. Jun 2022.

2021

T39 Visual learning in the open world. University of Oxford. Oxford, UK. Nov 2021.

T38 Visual learning in the open world. Google Brain. Toronto, Ontario, Canada. Nov 2021.

T37 Visual learning in the open world. Stanford University. Stanford, California, USA. Oct 2021.

T36 Steps towards making machine learning more natural. École Polytechnique Fédérale de Lausanne. Lausanne, Switzerland. Apr 2021.

T35 Steps towards making machine learning more natural. University of Michigan. Ann Arbor, Michigan, USA. Mar 2021.

T34 Steps towards making machine learning more natural. Université de Montréal. Montréal, Québec, Canada. Mar 2021.

T33 Steps towards making machine learning more natural. University of North Carolina, Chapel Hill. Chapel Hill, North Carolina, USA. Mar 2021.

- T32 Steps towards making machine learning more natural. University of Chicago. Chicago, Illinois, USA. Mar 2021.
- T31 Steps towards making machine learning more natural. University of British Columbia. Vancouver, British Columbia, Canada. Mar 2021.
- T30 Steps towards making machine learning more natural. University of Waterloo. Waterloo, Ontario, Canada. Mar 2021.
- T29 Steps towards making machine learning more natural. New York University. New York, New York, USA. Mar 2021.
- T28 Steps towards making machine learning more natural. University of Edinburgh. Edinburgh, UK. Mar 2021.
- T27 Steps towards making machine learning more natural. University of Maryland, College Park. College Park, Maryland, USA. Feb 2021.
- T26 A tutorial on few-shot learning and unsupervised representation learning. Vector Institute. Toronto, Ontario, Canada. Jan 2021.

2020

- T25 How can we apply few-shot learning? Vector Institute. Toronto, Ontario, Canada. Oct 2020.
- T24 Towards continual and compositional few-shot learning. Stanford University. Stanford, California, USA. Oct 2020.
- T23 Towards continual and compositional few-shot learning. Brown University. Providence, Rhode Island, USA. Sept 2020.
- T22 Towards continual and compositional few-shot learning. MIT. Cambridge, Massachusetts, USA. Sept 2020.
- T21 Towards continual and compositional few-shot learning. Mila. Montréal, Québec, Canada. Aug 2020.
- T20 Towards continual and compositional few-shot learning. Uber ATG. Toronto, Ontario, Canada. Aug 2020.
- T19 Wandering within a world: Online contextualized few-shot learning. Google Brain. Montréal, Québec, Canada. Aug 2020.
- T18 Wandering within a world: Online contextualized few-shot learning. ICML 2020 Lifelong Learning Workshop. Jul 2020.
- T17 Wandering within a world: Online contextualized few-shot learning. ICML 2020 Continual Learning Workshop. Jul 2020.

2019

- T16 Jointly learnable behavior and trajectory planner for autonomous driving. IROS 2019. Macau, China. Nov 2019.
- T15 Meta-learning for more human-like learning algorithms. Columbia University. New York, New York, USA. Oct 2019.

2018

- T14 Learning to reweight examples for robust deep learning. CIFAR deep learning and reinforcement learning summer school. Toronto, Ontario, Canada. Aug 2018.
- T13 Learning to reweight examples for robust deep learning. ICML 2018. Stockholm, Sweden. Jul 2018.
- T12 Meta-learning for weakly supervised learning. INRIA Grenoble - Rhône-Alpes. Grenoble, France. Jul 2018.

- T11 Meta-learning for weakly supervised learning. NEC Laboratories America. Princeton, New Jersey, USA. Jul 2018.
- T10 Meta-learning and learning to reweight examples. Max Planck Institute for Intelligent Systems. Tübingen, Germany. Jun 2018.
- T9 SBNNet: Sparse blocks network for fast inference. CVPR 2018. Salt Lake City, Utah, USA. Jun 2018.
- T8 SBNNet: Sparse blocks network for fast inference. Borealis AI Lab (RBC Research). Toronto, Ontario, Canada. Feb 2018.
- 2017**
- T7 Meta-learning for semi-supervised few-shot classification. Vector Institute. Toronto, Ontario, Canada. Nov 2017.
- T6 End-to-end instance segmentation with recurrent attention. CVPR 2017. Honolulu, Hawaii, USA. Jul 2017.
- T5 Sequence-to-sequence deep learning with recurrent attention. Queen’s University. Kingston, Ontario, Canada. May 2017.
- T4 Recurrent neural networks. CSC 2541: Sport Analytics Guest Lecture. University of Toronto. Toronto, Ontario, Canada. Jan 2017.
- 2016**
- T3 Deep dashboard tutorial. University of Guelph. Guelph, Ontario, Canada. Mar 2016.
- T2 Deep dashboard tutorial. University of Toronto. Toronto, Ontario, Canada. Feb 2016.
- 2015**
- T1 Exploring data and models for image question answering. ICML 2015 Deep Learning Workshop. Lille, France. Jul 2015.

STUDENT
SUPERVISION

Current Ph.D. Students:

- Chris Hoang 2023 - Present
- Jack (Yi Ao) Lu 2023 - Present
- Ryan Teehan 2022 - Present
- Ying Wang 2024 - Present
- Yanlai Yang 2022 - Present

Past Ph.D. Students:

- Alexander N. Wang 2022 - 2025

MEDIA COVERAGE

- The AI science separation. *The Wire China*. [[link](#)]. 2026/04/19.
- There’s a new ‘Reddit for AIs’, with no humans allowed. Is it proof they’re about to take over – or a sci-fi nothingburger? *The Independent*. [[link](#)]. 2026/02/05.
- What vision AI could be used for in the future. *Marketplace*. [[link](#)]. 2025/08/21.
- Could A.I. turn cities into economic losers? *CBS News New York*. [[link](#)]. 2025/04/10.
- Interview: Even experts don’t know how AI makes decisions. *MarketWatch*. [[link](#)]. 2024/02/15.
- Panel: Can AI be regulated and how? *CGTN*. [[link](#)]. 2023/06/16.
- Researchers Build AI That Builds AI. *Quanta Magazine*. [[link](#)]. 2022/01/25.