

Tongzhou Wang

✉ tongzhou@mit.edu | 🏠 tongzhouwang.info | 📄 Google Scholar | 📷 ssnl

EDUCATION

Massachusetts Institute of Technology

Ph.D. in Computer Science

2019 - 2024

- Advisors: Antonio Torralba, Phillip Isola
- Thesis: Intelligent Agents via Representation Learning

Massachusetts Institute of Technology

M.S. in Computer Science

2019 - 2022

- Advisors: Antonio Torralba, Phillip Isola
- Thesis: Geometric Properties of Learned Representations

University of California, Berkeley

B.A. in Computer Science and Statistics

2013 - 2017

- Advisors: Stuart J. Russell, Ren Ng, Alexei A. Efros

EXPERIENCES

Simons Institute for the Theory of Computing

Long-Term Visitor

2024

- Summer Cluster: AI, Psychology, and Neuroscience.

Facebook AI Research (FAIR)

Research Intern

2021

- Mentor: Yuandong Tian. Minimal world model for reinforcement learning. Paper published in ICML 2022.

Facebook AI Research (FAIR)

Software Engineer

2017 - 2019

- Built PyTorch. Worked on data pipelines, autograd, ML operators, etc.

RESEARCH INTERESTS

Learning world representations for generalist agents.

I am interested in learning structured representations that aggregate and select information about the world from various data sources, improve multi-task training, and enable autonomous adaptation to new tasks.

FEATURED PUBLICATIONS

(* indicates equal contribution)

The Platonic Representation Hypothesis

Minyoung Huh*, Brian Cheung*, **Tongzhou Wang***, Phillip Isola*

2024

International Conference on Machine Learning 2024 [ICML 2024 (Position Paper)]

Optimal Goal-Reaching Reinforcement Learning via Quasimetric Learning

Tongzhou Wang, Antonio Torralba, Phillip Isola, Amy Zhang

2023

International Conference on Machine Learning 2023 [ICML 2023]

Denoised MDPs: Learning World Models Better Than the World Itself

Tongzhou Wang, Simon S. Du, Antonio Torralba, Phillip Isola, Amy Zhang, Yuandong Tian

2022

International Conference on Machine Learning 2022 [ICML 2022]

Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere

Tongzhou Wang, Phillip Isola

2020

International Conference on Machine Learning 2020 [ICML 2020].

Dataset Distillation

Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, Alexei A. Efros

2018

Learning to See by Looking at Noise

Manel Baradad*, Jonas Wulff*, **Tongzhou Wang**, Phillip Isola, Antonio Torralba

2021

Advances in Neural Information Processing Systems 2021 [NeurIPS 2021]

INVITED TALKS

Reinforcement Learning as Representation Learning

Simons Institute

June 2024

UC Berkeley

May 2024

Structured Representations for Active Agents

Stanford Vision and Learning Lab (SVL), Stanford University

November 2023

Guest Lecture, USC

November 2023

Quasimetric Reinforcement Learning

Brown University

November 2023

AI Seminar, CMU

October 2023

Vector Institute

September 2023

Deep Learning: Classics and Trends (DLCT)

June 2023

Machine Learning Advances Symposium, MIT

May 2023

UT Austin

April 2023

Northeastern University

April 2023

Technical Talks on PyTorch Internals

PyTorch Developer Conference, San Francisco, CA, USA

October 2019

Global Mobile Internet Conference, Beijing, China

April 2018

MENTORING

Massachusetts Institute of Technology

Adrian Rodriguez Munoz (Ph.D. student)

Spring 2024 - PRESENT

Hyojin Bahng (Ph.D. student)

Summer 2023 - PRESENT

David X. Wu (B.S. & M.S. '22; now Ph.D. student at UC Berkeley)

Summer & Fall 2021

Jingwei Ma (B.S. & M.S. '21; now Ph.D. student at University of Washington)

2019 - 2022

Steven Liu (B.S. & M.S. '21; now at TwoSigma)

2019 - 2020

Carnegie Mellon University

George Cazenavette (M.S. '22; now Ph.D. student at MIT)

2021 - 2023

Summer Geometry Initiative (SGI)

Daniel Perazzo (master student at IMPA, Brazil)

Summer 2023 - PRESENT

Biruk Abere (B.S. student at University of Gondar, Ethiopia)

Summer 2023

Gabriele Dominici (master student at University of Cambridge, UK)

Summer 2023

Sana Arastehfar (master student at Queen's University, Canada)

Summer 2023

Sanowar Raihan (research assistant at Center for Computational & Data Sciences, Bangladesh)

Summer 2023

TEACHING

6.S898: Deep Learning, MIT

Fall 2022

Teaching Assistant (Co-Designed Curriculum and Assignments for 1st Undergraduate Offering)

Professional Development Course on Deep Learning, MIT

Summer 2019

Lab Session Instructor

Deep Learning Tutoring

Spring & Summer 2023

Volunteer Tutoring for a Data Science Professional in Boston, MA, USA

Deep Learning with PyTorch

Spring 2018

Tutorial and Lab Session Instructor (200-300 participants) at Global Mobile Internet Conference, Beijing, China

Middle-School Mathematics and English

Summer 2011

Volunteer Teaching for Low-Income Students in Northwestern China

SERVICES

Reviewer

ICML 2020-2024, NeurIPS 2020-2023, ICLR 2022, RLC 2024, CVPR 2021, TMLR, TPAMI, GCRL Workshop 2023.

Workshop Organizer

Goal-Conditioned Reinforcement Learning (GCRL) Workshop at NeurIPS 2023.
The First Dataset Distillation Challenge at ECCV 2024.

OPEN-SOURCE PROJECTS _____ (106k stars on GitHub combined over projects that I made significant contributions to)

PyTorch Core dev on data loading pipelines, CUDA/CPU kernels, ML ops, API design, autograd optimization, Python bindings, etc.	2017-2020
CycleGAN and pix2pix in PyTorch Maintainer for image-to-image translation toolkit	2018-PRESENT
torchreparam Built one of the first tools to reparametrize neural nets; now part of the large-scale SFDP training framework <code>fairscale</code>	2019-2020
torchqmet Built the first toolkit for learning quasimetrics with neural nets	2022-PRESENT

HONORS AND AWARDS _____

Meta Ph.D. Fellowship Finalist	2023
Outstanding Reviewer for ICML 2022	2022
Top Reviewer for ICML 2020	2020
Merrill Lynch Graduate Fellowship	2019
UC Berkeley High Distinction in General Scholarship	2017
Best Summer Social Practice of Shanghai for my volunteer teaching in northwestern China	2011

OTHER ENGINEERING EXPERIENCES _____

Airbnb, Inc. Intern on Machine Learning Infrastructure	2016
Facebook, Inc. Intern on Ads API Platform	2015
Grue, Inc. Co-Founder	2015
Sellegit, Inc. Intern on Mobile App Development	2014

PUBLICATIONS (COMPLETE LIST) _____ (* indicates equal contribution)

A Framework for Evaluating the Efficacy of Foundation Embedding Models in Healthcare Sonnet Xu, Haiwen Gui, Veronica Rotemberg, Tongzhou Wang , Yiqun Chen, Roxana Daneshjoui • medRxiv	2024
The Platonic Representation Hypothesis Minyoung Huh*, Brian Cheung*, Tongzhou Wang* , Phillip Isola* • International Conference on Machine Learning 2024 [ICML 2024 (Position Paper)] • Code Webpage arXiv	2024
Optimal Goal-Reaching Reinforcement Learning via Quasimetric Learning Tongzhou Wang , Antonio Torralba, Phillip Isola, Amy Zhang • International Conference on Machine Learning 2023 [ICML 2023]. • Code Webpage arXiv	2023
Generalizing Dataset Distillation via Deep Generative Prior George Cazenavette, Tongzhou Wang , Antonio Torralba, Alexei A. Efros, Jun-Yan Zhu • Conference on Computer Vision and Pattern Recognition 2023 [CVPR 2023]. • Code Webpage arXiv	2023
Deep Augmentation: Enhancing Self-Supervised Learning through Transformations in Higher Activation Space Rickard Br�uel-Gabrielsson, Tongzhou Wang , Manel Baradad, Justin Solomon • arXiv	2023

Steerable Equivariant Representation Learning

Sangnie Bhardwaj, Willie McClinton, [Tongzhou Wang](#), Guillaume Lajoie, Chen Sun, Phillip Isola, Dilip Krishnan

2023

- [arXiv](#)

Improved Representation of Asymmetrical Distances with Interval Quasimetric Embeddings

[Tongzhou Wang](#), Phillip Isola

2022

- Workshop on Symmetry and Geometry in Neural Representations at NeurIPS 2022 [[NeurReps Workshop at NeurIPS 2022](#)].
- [PyTorch Package for Quasimetric Learning](#) [Webpage](#) [OpenReview](#) [arXiv](#)

Procedural Image Programs for Representation Learning

Manel Baradad, Chun-Fu Chen, Jonas Wulff, [Tongzhou Wang](#), Rogerio Feris, Antonio Torralba, Phillip Isola

2022

- Advances in Neural Information Processing Systems 2022 [[NeurIPS 2022](#)].
- [Code & Datasets](#) [Webpage](#) [OpenReview](#) [arXiv](#)

Denoised MDPs: Learning World Models Better Than the World Itself

[Tongzhou Wang](#), Simon S. Du, Antonio Torralba, Phillip Isola, Amy Zhang, Yuandong Tian

2022

- International Conference on Machine Learning 2022 [[ICML 2022](#)].
- [Code](#) [Webpage](#) [arXiv](#)

On the Learning and Learnability of Quasimetrics

[Tongzhou Wang](#), Phillip Isola

2022

- International Conference on Learning Representations 2022 [[ICLR 2022](#)].
- [Code](#) [Webpage](#) [OpenReview](#) [arXiv](#)

Dataset Distillation by Matching Training Trajectories

George Cazenavette, [Tongzhou Wang](#), Antonio Torralba, Alexei A. Efros, Jun-Yan Zhu

2022

- Conference on Computer Vision and Pattern Recognition 2022 [[CVPR 2022](#)].
- [Code](#) [Webpage](#) [arXiv](#)

Wearable ImageNet: Synthesizing Tileable Textures via Dataset Distillation

George Cazenavette, [Tongzhou Wang](#), Antonio Torralba, Alexei A. Efros, Jun-Yan Zhu

2022

- 5th Workshop on Computer Vision for Fashion, Art, and Design at CVPR 2022 [[CVFAD Workshop at CVPR 2022](#)].
- [Code](#) [Webpage](#) [Paper](#)

Totems: Physical Objects for Verifying Visual Integrity

Jingwei Ma, Lucy Chai, Minyoung Huh, [Tongzhou Wang](#), Ser-Nam Lim, Phillip Isola, Antonio Torralba

2022

- European Conference on Computer Vision 2022 [[ECCV 2022](#)].
- [Code](#) [Webpage](#) [arXiv](#)

Learning to See by Looking at Noise

Manel Baradad*, Jonas Wulff*, [Tongzhou Wang](#), Phillip Isola, Antonio Torralba

2021

- Advances in Neural Information Processing Systems 2021 [[NeurIPS 2021](#)].
- [Code & Datasets](#) [Webpage](#) [arXiv](#)

Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere

[Tongzhou Wang](#), Phillip Isola

2020

- International Conference on Machine Learning 2020 [[ICML 2020](#)].
- [Code](#) [Webpage](#) [arXiv](#)

Rewriting a Deep Generative Model

David Bau, Steven Liu, [Tongzhou Wang](#), Jun-Yan Zhu, Antonio Torralba

2020

- European Conference on Computer Vision 2020 [[ECCV 2020](#)].
- [Code](#) [Webpage](#) [arXiv](#)

Diverse Image Generation via Self-Conditioned GANs

Steven Liu, [Tongzhou Wang](#), David Bau, Jun-Yan Zhu, Antonio Torralba

2020

- Conference on Computer Vision and Pattern Recognition 2020 [[CVPR 2020](#)].
- [Code](#) [Webpage](#) [arXiv](#)

Dataset Distillation

[Tongzhou Wang](#), Jun-Yan Zhu, Antonio Torralba, Alexei A. Efros

2018

- [Code](#) [Webpage](#) [arXiv](#)

Meta-Learning MCMC Proposals

Tongzhou Wang, Yi Wu, David A. Moore, Stuart J. Russell

2017

- Advances in Neural Information Processing Systems 2018 [[NeurIPS 2018](#)].
- Automatic Machine Learning Workshop at ICML 2017 (Oral) [[AutoML Workshop at ICML 2017 \(Oral\)](#)].
- [↗ arXiv](#)

Learning to Synthesize a 4D RGBD Light Field from a Single Image

Pratul Srinivasan, Tongzhou Wang, Ashwin Sreelal, Ravi Ramamoorthi, Ren Ng

2017

- International Conference on Computer Vision 2017 [[ICCV 2017](#)].
- [↗ Code arXiv](#)