

YUANDONG TIAN

RESEARCH INTEREST

Reinforcement Learning and Optimization, Representation Learning

WORK EXPERIENCE

- Jan.2015-present **Research Scientist and Senior Manager**, Meta AI (FAIR)
Research Lead in:
 AI-guided Optimization
 Self-supervised and Representation Learning
Lead Scientist and Engineer of Facebook Go engine.
 DarkForestGo (2015): Strong CNN-based Go AI before AlphaGo.
 ELF OpenGo (2018): Superhuman and best open source Go AI
 that beats professional Go players with 20-0.
 Replication of AlphaZero with 2k GPUs.
- Sep.2013-Dec.2014 **Researcher / Software Engineer**
Vision and Learning Group, Driverless Car Team (Waymo), Google X
Real-time object recognition for autonomous driving car.

EDUCATION

- Sep.2008-Sep.2013 **Ph.D** in Robotics
Carnegie Mellon University
- Sep.2005-Jun.2008 **Master** in Computer Science & Engineering
Shanghai Jiao Tong University
- Sep.2001-Jun.2005 **Bachelor** in Computer Science & Engineering
Shanghai Jiao Tong University

AWARDS AND SCHOLARSHIP

- Jul.2021 Recipient of ICML 2021 **Outstanding Paper Honorable Mentions.**
Dec.2013 Recipient of ICCV 2013 **Marr Prize Honorable Mentions.**

PAPER REVIEWS

- Reviewers CVPR, ICCV, ECCV, Eurographics, ICCP, ACM SIGCHI, AAAI
PAMI, IJCV, TIP, ICML, NeurIPS, ICLR, UAI
- Area Chair AAAI 2020, AISTats 2020-2022, NeurIPS 2021-2022

SKILLS

- Software C/C++, Python, LUA/torch, Matlab
Languages Mandarin Chinese (Native), English, Japanese

REINFORCEMENT LEARNING and OPTIMIZATION

- [1] Tongzhou Wang, Simon S Du, Antonio Torralba, Phillip Isola, Amy Zhang, **Yuandong Tian**, *Denoised MDPs: Learning World Models Better Than the World Itself*, ICML 2022
- [2] Yiyang Zhao, Linnan Wang, Kevin Yang, Tianjun Zhang, Tian Guo, **Yuandong Tian**, *Multi-objective Optimization by Learning Space Partitions*, ICLR 2022

- [3] Daochen Zha, Louis Feng, Bhargav Bhushanam, Dhruv Choudhary, Jade Nie, **Yuandong Tian**, Jay Chae, Yinbin Ma, Arun Kejariwal, Xia Hu, *AutoShard: Automated Embedding Table Sharding for Recommender Systems*, KDD 2022
- [4] Tianjun Zhang, Paria Rashidinejad, Jiantao Jiao, **Yuandong Tian**, Joseph Gonzalez, Stuart Russell, *MADE: Exploration via Maximizing Deviation from Explored Regions*, NeurIPS 2021
- [5] Tianjun Zhang, Huazhe Xu, Xiaolong Wang, Yi Wu, Kurt Keutzer, Joseph E. Gonzalez, **Yuandong Tian**, *Noveld: A simple yet effective exploration criterion*, NeurIPS 2021
- [6] Yiyang Zhao*, Linnan Wang*, **Yuandong Tian**, Rodrigo Fonseca, Tian Guo, *Few-shot Neural Architecture Search*, ICML 2021 (**Long oral**)
- [7] Hang Zhu, Varun Gupta, Satyaajeet Singh Ahuja, **Yuandong Tian**, Ying Zhang, Xin Jin, *Network Planning with Deep Reinforcement Learning*, SIGCOMM 2021
- [8] Tianjun Zhang, Huazhe Xu, Xiaolong Wang, Yi Wu, Kurt Keutzer, Joseph E Gonzalez, **Yuandong Tian**, *Multi-Agent Collaboration via Reward Attribution Decomposition*, arXiv 2020
- [9] Linnan Wang, Rodrigo Fonseca, **Yuandong Tian**, *Learning Search Space Partition for Black-box Optimization using Monte Carlo Tree Search*, NeurIPS 2020
- [10] **Yuandong Tian**, Qucheng Gong, Tina Jiang, *Joint Policy Search for Multi-agent Collaboration with Imperfect Information*, NeurIPS 2020
- [11] Hui Shi, Yang Zhang, Xinyun Chen, **Yuandong Tian**, Jishen Zhao, *Deep Symbolic Superoptimization Without Human Knowledge*, ICLR 2020.
- [12] Linnan Wang, Yiyang Zhao, Yuu Jinnai, **Yuandong Tian**, Rodrigo Fonseca, *AlphaX: Exploring Neural Architectures with Deep Neural Networks and Monte Carlo Tree Search*, AAAI 2020
- [13] Xinyun Chen, **Yuandong Tian**, *Learning to Perform Local Rewriting for Combinatorial Optimization*, NeurIPS 2019.
- [14] Hengyuan Hu*, Denis Yarats*, Qucheng Gong, **Yuandong Tian**, Mike Lewis, *Hierarchical Decision Making by Generating and Following Natural Language Instructions*, NeurIPS 2019.
- [15] **Yuandong Tian**, Jerry Ma*, Qucheng Gong*, Shubho Sengupta*, Zhuoyuan Chen, James Pinkerton, C. Lawrence Zitnick, *ELF OpenGo: An Analysis and Open Reimplementation of AlphaZero*, ICML 2019 (**Long oral**)
- [16] Yuping Luo, Harry Huazhe Xu, Yuanzhi Li, **Yuandong Tian**, Trevor Darrell, Tengyu Ma: *Algorithmic Framework for Model-based Deep Reinforcement Learning with Theoretical Guarantees*, ICLR 2019.
- [17] Tianmin Shu, **Yuandong Tian**: *M³RL: Mind-aware Multi-agent Management Reinforcement Learning*, ICLR 2019.
- [18] Yi Wu, Yuxin Wu, Aviv Tamar, Stuart Russell, Georgia Gkioxari, **Yuandong Tian**, *Bayesian Relational Memory for Semantic Visual Navigation*, ICCV 2019.
- [19] **Yuandong Tian**, Qucheng Gong, Wenling Shang, Yuxin Wu, C. Lawrence Zitnick, *ELF: An Extensive, Lightweight and Flexible Research Platform for Real-time Strategy Games*, NIPS 2017 (**Oral**).
- [20] Yuxin Wu, **Yuandong Tian**, *Training Agent for First-Person Shooter Game with Actor-Critic Curriculum Learning*, ICLR 2017.

[21] **Yuandong Tian**, Yan Zhu, *Better Computer Go Player with Neural Network and Long-term Prediction*, ICLR 2016.

UNDERSTANDING REPRESENTATION LEARNING

[1] **Yuandong Tian**, *Understanding the Role of Nonlinearity in Training Dynamics of Contrastive Learning*, arXiv 2022

[2] **Yuandong Tian**, *Understanding Deep Contrastive Learning via Coordinate-wise Optimization*, arXiv 2022

[3] Xiao Wang, Haoqi Fan, **Yuandong Tian**, Daisuke Kihara, Xinlei Chen, *On the Importance of Asymmetry for Siamese Representation Learning*, CVPR 2022

[4] Li Jing, Pascal Vincent, Yann LeCun, **Yuandong Tian**, *Understanding dimensional collapse in contrastive self-supervised learning*, ICLR 2022

[5] Xiang Wang, Xinlei Chen, Simon S Du, **Yuandong Tian**, *Towards demystifying representation learning with non-contrastive self-supervision*, arXiv, 2021

[6] **Yuandong Tian**, Xinlei Chen, Surya Ganguli, *Understanding Self-supervised Learning Dynamics without Contrastive Pairs*, ICML 2021 (**Outstanding Paper Honorable mention**)

[7] Zhuolin Yang*, Zhaoxi Chen, Tiffany Cai, Xinyun Chen, Bo Li, **Yuandong Tian***, *Understanding Robustness in Teacher-Student Setting: A New Perspective*, AISTATS 2021

[8] **Yuandong Tian**, Lantao Yu, Xinlei Chen, Surya Ganguli, *Understanding Self-supervised Learning with Dual Deep Networks*, arXiv, 2020

[9] **Yuandong Tian**, *Student Specialization in Deep ReLU Networks With Finite Width and Input Dimension*, ICML 2020

[10] Haonan Yu, Sergey Edunov, **Yuandong Tian**, Ari S Morcos, *Playing the lottery with rewards and multiple languages: lottery tickets in RL and NLP*, ICLR 2020.

[11] Ari Morcos, Haonan Yu, Michela Paganini, **Yuandong Tian**, *One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers*, NeurIPS 2019.

[12] **Yuandong Tian**, Tina Jiang, Qucheng Gong, Ari Morcos, *Luck Matters: Understanding Training Dynamics of Deep ReLU Networks*, arXiv 2019.

[13] Simon S. Du, Jason D. Lee, **Yuandong Tian**, Barnabas Poczos, Aarti Singh, *Gradient Descent Learns One-hidden-layer CNN: Don't be Afraid of Spurious Local Minima*, ICML 2018.

[14] Simon S. Du, Jason D. Lee, **Yuandong Tian**, *When is a Convolutional Filter Easy to Learn?* ICLR 2018.

[15] **Yuandong Tian**, *An Analytical Formula of Population Gradient for two-layered ReLU network and its Applications in Convergence and Critical Point Analysis*, ICML 2018.

NEURAL ARCHITECTURE OPTIMIZATION

[1] Chengyue Gong, Dilin Wang, Meng Li, Xinlei Chen, Zhicheng Yan, **Yuandong Tian**, Vikas Chandra, *NASViT: Neural Architecture Search for Efficient Vision Transformers with Gradient Conflict aware Supernet Training*, ICLR 2022

- [2] Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Bichen Wu, Zijian He, Zhen Wei, Kan Chen, **Yuandong Tian**, Matthew Yu, Peter Vajda, Joseph Gonzalez, *FBNetV3: Joint Architecture-Recipe Search using Predictor Pretraining*, CVPR 2021
- [3] Zhicheng Yan, Xiaoliang Dai, Peizhao Zhang, **Yuandong Tian**, Bichen Wu, Matt Feiszli, *FP-NAS: Fast Probabilistic Neural Architecture Search*, CVPR 2021
- [4] Cheng Fu, Hanxian Huang, Xinyun Chen, **Yuandong Tian**, Jishen Zhao, *Learn-to-Share: A Hardware-friendly Transfer Learning Framework Exploiting Computation and Parameter Sharing*, ICML 2021 (**Long oral**)
- [5] Linnan Wang, Saining Xie, Teng Li, Rodrigo Fonseca, **Yuandong Tian**, *Sample-Efficient Neural Architecture Search by Learning Action Space*, TPAMI 2021
- [6] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, **Yuandong Tian**, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, Peter Vajda, Joseph Gonzalez, *FBNetV2: Differentiable Neural Architecture Search for Spatial and Channel Dimensions*, CVPR 2020
- [7] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, **Yuandong Tian**, Peter Vajda, Yangqing Jia, Kurt Keutzer: *FBNet: Hardware-Aware Efficient ConvNet Design via Differentiable Neural Architecture Search*, CVPR 2019.

PROGRAM GENERATION and OPTIMIZATION

- [1] Chris Cummins, Bram Wasti, Jiadong Guo, Brandon Cui, Jason Ansel, Sahir Gomez, Somya Jain, Jia Liu, Olivier Teytaud, Benoit Steiner, **Yuandong Tian**, Hugh Leather, *CompilerGym: robust, performant compiler optimization environments for AI research*, CGO 2022 (**Outstanding Paper**)
- [2] Xinyun Chen, Dawn Song, **Yuandong Tian**, *Latent execution for neural program synthesis beyond domain-specific languages*, NeurIPS 2021
- [3] Cheng Fu, Kunlin Yang, Xinyun Chen, **Yuandong Tian**, Jishen Zhao, *N-Bref : A High-fidelity Decompiler Exploiting Programming Structures*, arXiv 2020
- [4] Cheng Fu, Huili Chen, Haolan Liu, Xinyun Chen, **Yuandong Tian**, Farinaz Koushanfar, Jishen Zhao, *Coda: An End-to-End Neural Program Decompiler*, NeurIPS 2019.

OTHER TOPICS

- [1] Qingquan Song, Dehua Cheng, Hanning Zhou, Jiyan Yang, **Yuandong Tian**, Xia Hu, *Towards Automated Neural Interaction Discovery for Click-Through Rate Prediction*, KDD 2020
- [2] Yan Zhu, **Yuandong Tian**, Dimitris Mexatas, Piotr Dollar, *Semantic Amodal Segmentation*, CVPR 2017.
- [3] Jiajun Wu, Tianfan Xue, Joseph J. Lim, **Yuandong Tian**, Joshua B. Tenenbaum, Antonio Torralba, William T. Freeman, *Single Image 3D Interpreter Network*, ECCV 2016 (Oral presentation).
- [4] Jiajun Wu, Tianfan Xue, Joseph J. Lim, **Yuandong Tian**, Joshua B. Tenenbaum, Antonio Torralba, William T. Freeman, *3D Interpreter Networks for Viewer-Centered Wireframe Modeling*, International Journal of Computer Vision (IJCV), Jul. 2015.
- [5] Bolei Zhou, **Yuandong Tian**, Sainbayar Sukhbaatar, Arthur Szlam, Rob Fergus, *Simple Baseline for Visual Question Answering*, arXiv, 2016.

PhD WORKS

- [1] **Yuandong Tian**, Srinivasa G. Narasimhan, *Theory and Practice of Hierarchical Data-driven Descent for Optimal Deformation Estimation*, International Journal of Computer Vision (IJCV), Jul. 2015.
- [2] **Yuandong Tian**, Srinivasa G. Narasimhan: *Hierarchical Data-Driven Descent for Efficient Optimal Deformation Estimation*, ICCV 2013 (**Marr Prize Honorable Mentions**).
- [3] Nan Li, **Yuandong Tian**, William W. Cohen, Ken Koedinger: *Integrating Perceptual Learning with External World Knowledge in a Simulated Student*, Proceedings of the 16th International Conference on Artificial Intelligence in Education, 2013.
- [4] **Yuandong Tian**, C. Lawrence Zitnick, Srinivasa G. Narasimhan: *Exploring the Spatial Hierarchy of Mixture Models for Human Pose Estimation*, ECCV 2012
- [5] **Yuandong Tian**, Jun Zhu: *Learning from Crowds in the Presence of Schools of Thought*, SIGKDD 2012.
- [6] **Yuandong Tian**, Srinivasa G. Narasimhan, Alan J. Vannevel: *Depth from Optical Turbulence*, CVPR 2012.
- [7] **Yuandong Tian**, Srinivasa G. Narasimhan, *Globally Optimal Estimation of Nonrigid Image Distortion*, International Journal of Computer Vision (IJCV), Nov. 2011.
- [8] Mohit Gupta, **Yuandong Tian**, Srinivasa G. Narasimhan, Li Zhang, *A Combined Theory of Defocused Illumination and Global Light Transport*, International Journal of Computer Vision (IJCV), Sep. 2011.
- [9] **Yuandong Tian**, Srinivasa G. Narasimhan: *Rectification and 3D reconstruction of Curved Document Images*, CVPR 2011, (**Oral Presentation**)
- [10] Dong Huang, **Yuandong Tian**, Fernando De la Torre, *Local Isomorphism to Solve the Pre-image Problem in Kernel Methods*, CVPR 2011.
- [11] **Yuandong Tian**, Srinivasa G. Narasimhan: *A Globally Optimal Data-Driven Approach for Image Distortion Estimation*, CVPR 2010, (**Oral Presentation**).
- [12] **Yuandong Tian**, Srinivasa G. Narasimhan: *Seeing through water: Image restoration using model-based tracking*, ICCV 2009.
- [13] Mohit Gupta, **Yuandong Tian**, Srinivasa G. Narasimhan, Li Zhang: *(De)Focusing on Global Light Transport for Active Scene Recovery*, CVPR 2009, (**Oral Presentation**).