

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:
 LLM inference/planning (2022-)
 AI-guided Optimization (2019-)
 Self-supervised and Representation Learning (2020-)
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

- Jul. 2023 *Outstanding Paper* in Sampling and Discrete Optimization (SODS) Workshop
ICML 2023
(Project lead)
- Apr. 2022 *Distinguished Paper Award* in CGO 2022
(Project Co-lead)
- Jul. 2021 *Outstanding Paper Honorable Mentions* in ICML 2021
(First author, Project Lead)
- Dec. 2013 *Marr Prize Honorable Mentions* in ICCV 2013
(First author, Project Lead)

SELECTED PAPERS

- [1] Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, Yuandong Tian, *GaLore: Memory-Efficient LLM Training by Gradient Low-Rank Projection*, ICML 2024 (*Oral*)[[link](#)] [[code](#)]
- [2] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, Mike Lewis, *Efficient Streaming Language Models with Attention Sinks*, ICLR 2024 [[link](#)] [[code](#)] [[MIT News](#)] [[Yannic Kilcher's video introduction](#)] [[VentureBeat](#)] [[Huggingface library](#)] [[Intel extension of Transformers](#)] [[MLC Chat](#)] [[Y-combinator](#)]

- [3] Yuandong Tian, Xinlei Chen, Surya Ganguli, *Understanding Self-supervised Learning Dynamics without Contrastive Pairs*, ICML 2021 (**Outstanding Paper Award Honorable Mention**)[[link](#)][[code](#)][[video](#)][[Slides](#)][[Blogpost](#)][[Independent Reproduction](#)]
- [4] Aaron Ferber, Taoan Huang, Daochen Zha, Martin Schubert, Benoit Steiner, Bistra Dilkina, Yuandong Tian, *SurCo: Learning Linear Surrogates For Combinatorial Nonlinear Optimization Problems*, ICML 2023 (**Outstanding paper in Sampling and Optimization in Discrete Space (SODS) Workshop**)[[link](#)][[code](#)][[Slides](#)]
- [5] Yuandong Tian, Jerry Ma*, Qucheng Gong*, Shubho Sengupta*, Zhuoyuan Chen, James Pinkerton, Larry Zitnick, *ELF OpenGo: An Analysis and Open Reimplementation of AlphaZero*, ICML 2019 (**Long Oral**)[[link](#)][[code](#)][[website](#)][[pretrained model and game records](#)][[Blogpost](#)][[Talk](#)][[Forbes](#)][[TechCrunch](#)][[reddit](#)]
- [6] Xinyun Chen, Yuandong Tian, *Learning to Perform Local Rewriting for Combinatorial Optimization*, NeurIPS 2019 [[link](#)][[code](#)]
- [7] Yuandong Tian, Yan Zhu, *Better Computer Go Player with Neural Network and Long-term Prediction*, ICLR 2016 [[link](#)][[code](#)][[pretrained model](#)][[mit tech review](#)][[wired](#)]

LLM EFFICIENCY

- [1] Hanshi Sun, Zhuoming Chen, Xinyu Yang, Yuandong Tian, Beidi Chen, *TriForce: Lossless Acceleration of Long Sequence Generation with Hierarchical Speculative Decoding*, COLM 2024 [[link](#)][[code](#)]
- [2] Zechun Liu, Changsheng Zhao, Forrest Iandola, Chen Lai, Yuandong Tian, Igor Fedorov, Yunyang Xiong, Ernie Chang, Yangyang Shi, Raghuraman Krishnamoorthi, Liangzhen Lai, Vikas Chandra, *MobileLLM: Optimizing Sub-billion Parameter Language Models for On-Device Use Cases*, ICML 2024 [[link](#)]
- [3] Ruisi Cai, Yuandong Tian, Zhangyang Wang, Beidi Chen, *LoCoCo: Dropping In Convolutions for Long Context Compression*, ICML 2024 [[link](#)]
- [4] Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, Yuandong Tian, *GaLore: Memory-Efficient LLM Training by Gradient Low-Rank Projection*, ICML 2024 (**Oral**)[[link](#)][[code](#)]
- [5] Kevin Yang, Dan Klein, Asli Celikyilmaz, Nanyun Peng, Yuandong Tian, *RLCD: Reinforcement Learning from Contrast Distillation for Language Model Alignment*, ICLR 2024 [[link](#)][[code](#)]
- [6] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, Mike Lewis, *Efficient Streaming Language Models with Attention Sinks*, ICLR 2024 [[link](#)][[code](#)][[MIT News](#)][[Yannic Kilcher's video introduction](#)][[VentureBeat](#)][[Huggingface library](#)][[Intel extension of Transformers](#)][[MLC Chat](#)][[Y-combinator](#)]
- [7] Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Re, Clark Barrett, Zhangyang Wang, Beidi Chen, *H2O: Heavy-Hitter Oracle for Efficient Generative Inference of Large Language Models*, NeurIPS 2023 [[link](#)][[code](#)]
- [8] Shouyuan Chen, Sherman Wong, Liangjian Chen, Yuandong Tian, *Extending Context Window of Large Language Models via Positional Interpolation*, arXiv 2023 [[link](#)]
- [9] Zichang Liu, Jue Wang, Tri Dao, Tianyi Zhou, Binhang Yuan, Zhao Song, Anshumali Shrivastava, Ce Zhang, Yuandong Tian, Christopher Re, Beidi Chen, *Deja Vu: Contextual Sparsity for Efficient LLMs at Inference Time*, ICML 2023 (**Oral**)[[link](#)]

LLM AGENT

- [1] Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao, Yu Su, *TravelPlanner: A Benchmark for Real-World Planning with Language Agents*, ICML 2024 (*Spotlight*)[\[link\]](#) [\[code\]](#)
- [2] Hanlin Zhu*, Andrew Cohen*, Danqing Wang, Kevin Yang, Xiaomeng Yang, Jiantao Jiao, Yuandong Tian, *End-to-end Story Plot Generator*, arXiv 2023 [\[link\]](#)
- [3] Danqing Wang, Kevin Yang, Hanlin Zhu, Xiaomeng Yang, Andrew Cohen, Lei Li, Yuandong Tian, *Learning Personalized Story Evaluation*, arXiv 2023 [\[link\]](#) [\[code\]](#)
- [4] Kevin Yang, Dan Klein, Nanyun Peng, Yuandong Tian, *DOC: Improving Long Story Coherence With Detailed Outline Control*, ACL 2023 [\[link\]](#) [\[code\]](#)
- [5] Kevin Yang, Yuandong Tian, Nanyun Peng, Dan Klein, *Re3: Generating Longer Stories With Recursive Reprompting and Revision*, EMNLP 2022 [\[link\]](#) [\[code\]](#)

REINFORCEMENT LEARNING AND SEARCH

- [1] Lucas Lehnert, Sainbayar Sukhbaatar, Paul Mcvay, Michael Rabbat, Yuandong Tian, *Beyond A*: Better Planning with Transformers via Search Dynamics Bootstrapping*, COLM 2024 [\[link\]](#)
- [2] Zhengyao Jiang*, Yingchen Xu*, Nolan Wagener, Yicheng Luo, Michael Janner, Edward Grefenstette, Tim Rocktaschel, Yuandong Tian, *H-GAP: Humanoid Control with a Generalist Planner*, ICLR 2024 (*Spotlight*)[\[link\]](#) [\[website\]](#)
- [3] Zhengyao Jiang, Tianjun Zhang, Michael Janner, Yueying Li, Tim Rocktaschel, Edward Grefenstette, Yuandong Tian, *Efficient Planning in a Compact Latent Action Space*, ICLR 2023 [\[link\]](#) [\[code\]](#) [\[website\]](#)
- [4] 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 [\[link\]](#) [\[code\]](#) [\[website\]](#)
- [5] Tianjun Zhang*, Paria Rashidinejad*, Jiantao Jiao, Yuandong Tian, Joseph Gonzalez, Stuart Russell, *MADE: Exploration via Maximizing Deviation from Explored Regions*, NeurIPS 2021 [\[link\]](#) [\[code\]](#)
- [6] Tianjun Zhang, Huazhe Xu, Xiaolong Wang, Yi Wu, Kurt Keutzer, Joseph E. Gonzalez, Yuandong Tian, *NovelD: A Simple yet Effective Exploration Criterion*, NeurIPS 2021 [\[link\]](#) [\[code\]](#) [\[video\]](#)
- [7] Yuandong Tian, Qucheng Gong, Tina Jiang, *Joint Policy Search for Multi-agent Collaboration with Imperfect Information*, NeurIPS 2020 [\[link\]](#) [\[code\]](#) [\[video\]](#)
- [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 [\[link\]](#) [\[code\]](#) [\[video\]](#) [\[website\]](#)
- [9] Tianmin Shu, Yuandong Tian, *M³RL: Mind-aware Multi-agent Management Reinforcement Learning*, ICLR 2019 [\[link\]](#) [\[code\]](#)
- [10] Yuping Luo*, Huazhe Xu*, Yuanzhi Li, Yuandong Tian, Trevor Darrell, Tengyu Ma, *Algorithmic Framework for Model-based Deep Reinforcement Learning with Theoretical Guarantees*, ICLR 2019 [\[link\]](#) [\[code\]](#)
- [11] Yuandong Tian, Jerry Ma*, Qucheng Gong*, Shubho Sengupta*, Zhuoyuan Chen, James Pinker-

ton, Larry Zitnick, *ELF OpenGo: An Analysis and Open Reimplementation of AlphaZero*, ICML 2019 (**Long Oral**)[\[link\]](#) [\[code\]](#) [\[website\]](#) [\[pretrained model and game records\]](#) [\[Blogpost\]](#) [\[Talk\]](#) [\[Forbes\]](#) [\[TechCrunch\]](#) [\[reddit\]](#)

[12] Hengyuan Hu*, Denis Yarats*, Qucheng Gong, Yuandong Tian, Mike Lewis, *Hierarchical Decision Making by Generating and Following Natural Language Instructions*, NeurIPS 2019 [\[link\]](#) [\[code\]](#)

[13] Yi Wu, Yuxin Wu, Georgia Gkioxari, Yuandong Tian, *Building Generalizable Agents with a Realistic and Rich 3D Environment*, ICLR-Workshop 2018 [\[link\]](#) [\[code\]](#)

[14] Yuxin Wu, Yuandong Tian, *Training Agent for First-Person Shooter Game with Actor-Critic Curriculum Learning*, ICLR 2017 [\[link\]](#)

[15] Yuandong Tian, Qucheng Gong, Wenling Shang, Yuxin Wu, Larry Zitnick, *ELF: An Extensive, Lightweight and Flexible Research Platform for Real-time Strategy Games*, NeurIPS 2017 (**Oral**)[\[link\]](#) [\[code\]](#) [\[video\]](#)

[16] Yuandong Tian, Yan Zhu, *Better Computer Go Player with Neural Network and Long-term Prediction*, ICLR 2016 [\[link\]](#) [\[code\]](#) [\[pretrained model\]](#) [\[mit tech review\]](#) [\[wired\]](#)

UNDERSTANDING NEURAL NETWORKS

[1] Yuandong Tian, Yiping Wang, Zhenyu Zhang, Beidi Chen, Simon Du, *JoMA: Demystifying Multilayer Transformers via JOint Dynamics of MLP and Attention*, ICLR 2024 [\[link\]](#) [\[5min ICLR talk\]](#) [\[5min ICLR slides\]](#)

[2] Yuandong Tian, Yiping Wang, Beidi Chen, Simon Du, *Scan and Snap: Understanding Training Dynamics and Token Composition in 1-layer Transformer*, NeurIPS 2023 [\[link\]](#) [\[talk\]](#) [\[slides\]](#) [\[5min NeurIPS talk\]](#) [\[5min NeurIPS slides\]](#) [\[poster\]](#)

[3] Randall Balestriero, Mark Ibrahim, Vlad Sobal, Ari Morcos, Shashank Shekhar, Tom Goldstein, Florian Bordes, Adrien Bardes, Gregoire Mialon, Yuandong Tian, Avi Schwarzschild, Andrew Gordon Wilson, Jonas Geiping, Quentin Garrido, Pierre Fernandez, Amir Bar, Hamed Pirsiavash, Yann LeCun, Micah Goldblum, *A Cookbook of Self-Supervised Learning*, arXiv 2023 [\[link\]](#)

[4] Yuandong Tian, *Understanding the Role of Nonlinearity in Training Dynamics of Contrastive Learning*, ICLR 2023 [\[link\]](#) [\[code\]](#) [\[workshop version\]](#) [\[workshop poster\]](#) [\[5min talk\]](#)

[5] Yuandong Tian, *Understanding Deep Contrastive Learning via Coordinate-wise Optimization*, NeurIPS 2022 (**Oral**)[\[link\]](#) [\[code\]](#) [\[video\]](#) [\[5min talk slides\]](#) [\[poster\]](#)

[6] Xiao Wang*, Haoqi Fan*, Yuandong Tian, Daisuke Kihara, Xinlei Chen, *On the Importance of Asymmetry for Siamese Representation Learning*, CVPR 2022 [\[link\]](#) [\[code\]](#)

[7] Li Jing, Pascal Vincent, Yann LeCun, Yuandong Tian, *Understanding Dimensional Collapse in Contrastive Self-supervised Learning*, ICLR 2022 [\[link\]](#) [\[code\]](#)

[8] Xiang Wang, Xinlei Chen, Simon S Du, Yuandong Tian, *Towards demystifying representation learning with non-contrastive self-supervision*, arXiv 2021 [\[link\]](#)

[9] Hui Shi, Sicun Gao, Yuandong Tian, Xinyun Chen, Jishen Zhao, *Learning Bounded Context-Free-Grammar via LSTM and the Transformer: Difference and Explanations*, AAAI 2022 [\[link\]](#) [\[code\]](#)

[10] Yuandong Tian, Xinlei Chen, Surya Ganguli, *Understanding Self-supervised Learning Dynamics without Contrastive Pairs*, ICML 2021 (**Outstanding Paper Award Honorable Mention**)[\[link\]](#)

[\[code\]](#) [\[video\]](#) [\[Slides\]](#) [\[Blogpost\]](#) [\[Independent Reproduction\]](#)

[11] Zhuolin Yang*, Zhaoxi Chen, Tiffany Cai, Xinyun Chen, Bo Li, Yuandong Tian*, *Understanding Robustness in Teacher-Student Setting: A New Perspective*, AISTATS 2021 [\[link\]](#) [\[Slides\]](#)

[12] 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 [\[link\]](#)

[13] Yuandong Tian, *Student Specialization in Deep ReLU Networks With Finite Width and Input Dimension*, ICML 2020 [\[link\]](#) [\[code\]](#)

[14] Yuandong Tian, Lantao Yu, Xinlei Chen, Surya Ganguli, *Understanding Self-supervised Learning with Dual Deep Networks*, arXiv 2020 [\[link\]](#) [\[code\]](#) [\[video\]](#)

[15] Ari S. Morcos, Haonan Yu, Michela Paganini, Yuandong Tian, *One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers*, NeurIPS 2019 [\[link\]](#)

[16] Yuandong Tian, Tina Jiang, Qucheng Gong, Ari Morcos, *Luck Matters: Understanding Training Dynamics of Deep ReLU Networks*, ICML-workshop 2019 [\[link\]](#) [\[code\]](#) [\[Poster\]](#)

[17] Simon S. Du, Jason D. Lee, Yuandong Tian, *When is a Convolutional Filter Easy To Learn?*, ICLR 2018 [\[link\]](#)

[18] Simon S. Du, Jason D. Lee, Yuandong Tian, Barnabas Póczos, Aarti Singh, *Gradient Descent Learns One-hidden-layer CNN: Don't be Afraid of Spurious Local Minima*, ICML 2018 (**Long Oral**)[\[link\]](#)

[19] Yuandong Tian, *A Theoretical Framework for Deep Locally Connected ReLU Network*, arxiv 2018 [\[link\]](#) [\[Poster\]](#)

[20] Yuandong Tian, *An Analytical Formula of Population Gradient for two-layered ReLU network and its Applications in Convergence and Critical Point Analysis*, ICML 2017 [\[link\]](#) [\[code\]](#)

ML-GUIDED OPTIMIZATION

[1] Aaron Ferber, Arman Zharmagambetov, Taoan Huang, Bistra Dilkina, Yuandong Tian, *GenCO: Generating Diverse Solutions to Design Problems with Combinatorial Nature*, ICML 2024 [\[link\]](#)

[2] Taoan Huang, Aaron M Ferber, Arman Zharmagambetov, Yuandong Tian, Bistra Dilkina, *Contrastive Predict-and-Search for Mixed Integer Linear Programs*, ICML 2024

[3] Anselm Paulus*, Arman Zharmagambetov*, Chuan Guo, Brandon Amos**, Yuandong Tian**, *AdvPrompter: Fast Adaptive Adversarial Prompting for LLMs*, arXiv 2024 [\[link\]](#) [\[code\]](#)

[4] Arman Zharmagambetov, Brandon Amos, Aaron Ferber, Taoan Huang, Bistra Dilkina, Yuandong Tian, *Landscape Surrogate: Learning Decision Losses for Mathematical Optimization Under Partial Information*, NeurIPS 2023 [\[link\]](#) [\[code\]](#)

[5] Aaron Ferber, Taoan Huang, Daochen Zha, Martin Schubert, Benoit Steiner, Bistra Dilkina, Yuandong Tian, *SurCo: Learning Linear Surrogates For Combinatorial Nonlinear Optimization Problems*, ICML 2023 (**Outstanding paper in Sampling and Optimization in Discrete Space (SODS) Workshop**)[\[link\]](#) [\[code\]](#) [\[Slides\]](#)

[6] Taoan Huang, Aaron Ferber, Yuandong Tian, Bistra Dilkina, Benoit Steiner, *Searching Large Neighborhoods for Integer Linear Programs with Contrastive Learning*, ICML 2023 [\[link\]](#)

[7] Taoan Huang, Aaron Ferber, Yuandong Tian, Bistra Dilkina, Benoit Steiner, *Local Branching*

Relaxation Heuristics for Integer Linear Programs, CPAIOR 2023 [\[link\]](#)

[8] Andrew Cohen*, Weiping Dou, Jiang Zhu, Slawomir Koziel, Peter Renner, Jan-Ove Mattsson, Xiaomeng Yang, Beidi Chen, Kevin Stone, Yuandong Tian*, *Modeling Scattering Coefficients using Self-Attentive Complex Polynomials with Image-based Representation*, arXiv 2023 [\[link\]](#)

[9] Yiyang Zhao, Linnan Wang, Kevin Yang, Tianjun Zhang, Tian Guo, Yuandong Tian, *Multi-objective Optimization by Learning Space Partitions*, ICLR 2022 [\[link\]](#) [\[code\]](#)

[10] 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 [\[link\]](#) [\[code\]](#)

[11] Linnan Wang, Saining Xie, Teng Li, Rodrigo Fonseca, Yuandong Tian, *Sample-Efficient Neural Architecture Search by Learning Action Space*, T-PAMI 2021 [\[link\]](#) [\[code\]](#)

[12] Kevin Yang*, Tianjun Zhang*, Chris Cummins, Brandon Cui, Benoit Steiner, Linnan Wang, Joseph E. Gonzalez, Dan Klein, Yuandong Tian, *Learning Space Partitions for Path Planning*, NeurIPS 2021 [\[link\]](#) [\[code\]](#)

[13] Yiyang Zhao*, Linnan Wang*, Yuandong Tian, Rodrigo Fonseca, Tian Guo, *Few-shot Neural Architecture Search*, ICML 2021 (**Long Oral**)[\[link\]](#) [\[code\]](#) [\[Blogpost\]](#)

[14] Hang Zhu (JHU), Varun Gupta, Satyajeet Singh Ahuja, Yuandong Tian, Ying Zhang, Xin Jin, *Network Planning with Deep Reinforcement Learning*, SIGCOMM 2021 [\[link\]](#) [\[code\]](#)

[15] Zhicheng Yan, Xiaoliang Dai, Peizhao Zhang, Yuandong Tian, Bichen Wu, Matt Feiszli, *FPNAS: Fast Probabilistic Neural Architecture Search*, CVPR 2021 [\[link\]](#)

[16] 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 [\[link\]](#)

[17] 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 [\[link\]](#) [\[code\]](#)

[18] Qingquan Song, Dehua Cheng, Hanning Zhou, Jiyan Yang, Yuandong Tian, Xia Hu, *Towards Automated Neural Interaction Discovery for Click-Through Rate Prediction*, KDD 2020 [\[link\]](#)

[19] Linnan Wang, Rodrigo Fonseca, Yuandong Tian, *Learning Search Space Partition for Black-box Optimization using Monte Carlo Tree Search*, NeurIPS 2020 [\[link\]](#) [\[code\]](#)

[20] Xinyun Chen, Yuandong Tian, *Learning to Perform Local Rewriting for Combinatorial Optimization*, NeurIPS 2019 [\[link\]](#) [\[code\]](#)

[21] 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 [\[link\]](#) [\[code\]](#)

[22] Hongzi Mao, Shannon Chen, Drew Dimmery, Shaun Singh, Drew Blaisdell, Yuandong Tian, Mohammad Alizadeh, Eytan Bakshy, *Real-world video adaptation with reinforcement learning*, ICML-Workshop 2019 [\[link\]](#)

ML+SYS

[1] Youwei Liang*, Kevin Stone*, Ali Shameli, Chris Cummins, Mostafa Elhoushi, Jiadong Guo,

- Benoit Steiner, Xiaomeng Yang, Pengtao Xie, Hugh Leather, **Yuandong Tian**, *Learning Compiler Pass Orders using Coreset and Normalized Value Prediction*, ICML 2023 [[link](#)]
- [2] Daochen Zha, Louis Feng, Liang Luo, Bhargav Bhushanam, Zirui Liu, Yusuo Hu, Jade Nie, Yuzhen Huang, **Yuandong Tian**, Arun Kejariwal, Xia Hu, *Pre-train and Search: Efficient Embedding Table Sharding with Pre-trained Neural Cost Models*, MLSys 2023 [[link](#)] [[code](#)]
- [3] Jiaxun Cui, Xiaomeng Yang*, Geunbae Lee*, Mulong Luo*, Peter Stone, Hsien-Hsin S. Lee, Benjamin Lee, G. Edward Suh, Wenjie Xiong**, **Yuandong Tian****, *MACTA: A Multi-agent Reinforcement Learning Approach for Cache Timing Attacks and Detection*, ICLR 2023 [[link](#)]
- [4] Mulong Luo*, Wenjie Xiong*, Geunbae Lee, Yueying Li, Xiaomeng Yang, Amy Zhang, **Yuandong Tian**, Hsien Hsin S Lee, G Edward Suh, *AutoCAT: Reinforcement Learning for Automated Exploration of Cache Timing-Channel Attacks*, HPCA 2023 [[link](#)] [[code](#)]
- [5] Daochen Zha, Louis Feng, Qiaoyu Tan, Zirui Liu, Kwei-Herng Lai, Bhargav Bhushanam, **Yuandong Tian**, Arun Kejariwal, Xia Hu, *DreamShard: Generalizable Embedding Table Placement for Recommender Systems*, NeurIPS 2022 [[link](#)] [[code](#)]
- [6] 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 [[link](#)] [[code](#)]
- [7] 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**)[[link](#)] [[code](#)]
- [8] Xinyun Chen, Dawn Song, **Yuandong Tian**, *Latent Execution for Neural Program Synthesis Beyond Domain-Specific Languages*, NeurIPS 2021 [[link](#)] [[code](#)]
- [9] Cheng Fu, Hanxian Huang, Xinyun Chen, **Yuandong Tian**, Jishen Zhao (UCSD), *Learn-to-Share: A Hardware-friendly Transfer Learning Framework Exploiting Computation and Parameter Sharing*, ICML 2021 (**Long Oral**)[[link](#)]
- [10] Hui Shi, Yang Zhang, Xinyun Chen, **Yuandong Tian**, Jishen Zhao, *Deep Symbolic Superoptimization Without Human Knowledge*, ICLR 2020 [[link](#)] [[code](#)]
- [11] 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 [[link](#)]
- [12] Cheng Fu, Kunlin Yang, Xinyun Chen, **Yuandong Tian**, Jishen Zhao, *N-Bref: A High-fidelity Decompiler Exploiting Programming Structures*, arxiv 2020 [[link](#)] [[code](#)] [[Blogpost](#)]
- [13] Cheng Fu, Huili Chen, Haolan Liu, Xinyun Chen, **Yuandong Tian**, Farinaz Koushanfar, Jishen Zhao, *Coda: An End-to-End Neural Program Decompiler*, NeurIPS 2019 [[link](#)]

OTHER RESEARCH TOPICS

- [1] Yan Zhu, **Yuandong Tian**, Dimitris Mexatas, Piotr Dollár, *Semantic Amodal Segmentation*, CVPR 2017 [[link](#)]
- [2] Bolei Zhou, **Yuandong Tian**, Sainbayar Sukhbaatar, Arthur Szlam, Rob Fergus, *Simple Baseline for Visual Question Answering*, arxiv 2016 [[link](#)] [[code](#)]
- [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**)[[link](#)]

PHD WORK

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