

# TIANJIAO DING

[Homepage: ding.tj] ◊ [Google Scholar] ◊ [tjding@seas.upenn.edu] ◊ +1 (323) 828 8830 ◊ Philadelphia, PA, USA

## RESEARCH INTERESTS

---

I am interested in improving theoretical foundations of AI and developing efficient and trustworthy methods for AI. In particular, my research studies the parsimonious structures in data such as low-/high-dimensional linear subspaces, non-linear manifolds, and compositional semantics, and exploits them for emerging applications in vision and language.

- **Machine Learning:** flow matching, manifold clustering, high-rank subspace learning, low-dimensional representations
- **Computer Vision:** image clustering, semantic image generation, geometric computer vision, 3D scene analysis and synthesis
- **Language Modeling:** language model alignment, linear-time sequence modeling, self-supervised learning

## EDUCATION

---

**Ph.D. in Computer and Information Science** 2026 (Expected)  
University of Pennsylvania Philadelphia, USA

Advisor: [René Vidal] Committee: [Kostas Daniilidis], [Yi Ma], [Chris Callison-Burch], [Jeremias Sulam], [Hamed Hassani]  
Penn AI Fellow, CPAL Rising Stars

**M.S.E. in Applied Mathematics and Statistics** 2023  
Johns Hopkins University Baltimore, USA  
GPA: 3.98/4.00

**B.E. in Computer Science** 2018  
ShanghaiTech University Shanghai, China

Advisor: [Manolis C. Tsakiris]

Graduation with honor, Academic Excellence Scholarship 2014 & 2016

## AWARDS AND HONORS

---

**Penn AI Fellow**, University of Pennsylvania 2026

**Rising Stars Award**, Conference on Parsimony and Learning 2026

**Spotlight Paper (85/11603  $\approx$  0.7% acceptance rate)**, International Conference on Learning Representations 2025

**Schmidt Science Fellows Institutional Nominee**, 1 of 7 PhD students nominated by University of Pennsylvania 2025

**Winner for GRASP SFI internal Speakers**, University of Pennsylvania 2025

**Travel Grant Award**, Workshop on Seeking Low-Dimensionality in Deep Neural Networks 2022

**Graduation with Honor**, ShanghaiTech University 2018

**Academic Excellence Fellowship**  $\times 2$ , ShanghaiTech University 2014, 2016

**Outstanding Industrial Practice at National Instruments**, ShanghaiTech University 2016

## PUBLICATIONS

---

### Selected:

1. Token Statistics Transformer: Linear-Time Attention via Variational Rate Reduction [pdf] [ICLR 2025]  
Spotlight presentation, 85/11603  $\approx$  0.7% acceptance rate  
Ziyang Wu, **TD**, Yifu Lu, Druv Pai, Jingyuan Zhang, Weida Wang, Yaodong Yu, Yi Ma, and Benjamin D. Haeffele
2. PaCE: Parsimonious Concept Engineering for Large Language Models [pdf] [code] [video] [NeurIPS 2024]  
Jinqi Luo\*, **TD**\*, Kwan Ho Ryan Chan, Darshan Thaker, Aditya Chattopadhyay, Chris Callison-Burch, and René Vidal
3. Geometric Analysis of Nonlinear Manifold Clustering [pdf] [NeurIPS 2024]  
Top 116/3648  $\approx$  3% of accepted posters by reviewer rating  
Nimita Shinde\*, **TD**\*, Daniel P. Robinson, and René Vidal
4. Image Clustering via the Principle of Rate Reduction in the Age of Pretrained Models [pdf] [code] [ICLR 2024]

Tianzhe Chu\*, Shengbang Tong\*, **TD**\*, Xili Dai, Benjamin D. Haeffele, René Vidal, and Yi Ma

5. HARD: Hyperplane ARrangement Descent [pdf] [code] [CPAL 2024]  
**Oral presentation**  
**TD**\*, Liangzu Peng\*, and René Vidal
6. Unsupervised Manifold Linearizing and Clustering [pdf] [ICCV 2023]  
**CPAL spotlight**  
**TD**, Shengbang Tong, Kwan Ho Ryan Chan, Xili Dai, Yi Ma, and Benjamin D. Haeffele
7. Understanding Doubly Stochastic Clustering [pdf] [ICML 2022]  
**Oral presentation**  
**TD**, Derek Lim, René Vidal, and Benjamin D. Haeffele
8. Robust Homography Estimation via Dual Principal Component Pursuit [pdf] [code] [video] [CVPR 2020]  
**TD**, Yunchen Yang, Zhihui Zhu, Daniel P. Robinson, René Vidal, Laurent Kneip, and Manolis C. Tsakiris
9. Noisy Dual Principal Component Pursuit [pdf] [code] [ICML 2019]  
**Oral presentation**  
Tianyu Ding\*, Zhihui Zhu\*, **TD**, Yunchen Yang, Daniel P. Robinson, Manolis C. Tsakiris, and René Vidal
10. Learning to Parse Wireframes in Images of Man-Made Environments [pdf] [code] [CVPR 2018]  
**Over 258 citations as of Nov 2025**  
Kun Huang, Yifan Wang, Zihan Zhou, **TD**, Shenghua Gao, and Yi Ma

#### Other:

11. An Overview of Low-Rank Structures in the Training and Adaptation of Large Models [pdf] [arXiv 2025]  
Laura Balzano, **TD**, Benjamin D. Haeffele, Soo Min Kwon, Qing Qu, Peng Wang, Zhangyang Wang, and Can Yaras
12. Voyaging into Perpetual Dynamic Scenes from a Single View [pdf] [ICCV 2025]  
Fengrui Tian, **TD**, Jinqi Luo, Hancheng Min, and René Vidal
13. Concept Lancet: Representation Decomposition and Transplant for Diffusion-Based Image [pdf] [CVPR 2025]  
Jinqi Luo, **TD**, Kwan Ho Ryan Chan, Hancheng Min, Chris Callison-Burch, and René Vidal
14. Outlier-Robust Orthogonal Regression on Manifolds [OpenReview 2023]  
**TD**\*, Liangzu Peng\*, and René Vidal
15. Efficient Maximal Coding Rate Reduction by Variational Forms [pdf] [CVPR 2022]  
Christina Baek, Ziyang Wu, Kwan Ho Ryan Chan, **TD**, Yi Ma, and Benjamin D. Haeffele
16. Boosting RANSAC via Dual Principal Component Pursuit [arXiv 2021]  
Yunchen Yang, Xinyue Zhang, **TD**, Daniel P. Robinson, René Vidal, Manolis C. Tsakiris

## EMPLOYMENT

---

**Research Scientist Intern**, Agentic AI Team at Amazon, New York City, USA May - Aug. 2025

- Mentors: [Matthew Trager], [Aditya Chattopadhyay], [Elman Mansimov] Director: [Stefano Soatto]
- Project details subject to confidentiality agreement.

**Research Assistant**, Tsakiris Lab at ShanghaiTech, Shanghai, China 2017 - 2020

- Advisor: [Manolis C. Tsakiris] Collaborators: [Laurent Kneip], [René Vidal]
- Conducted detailed study of globally optimal subspace learning methods for outlier-robust 3D vision problems, e.g., 3D surface estimation [9], 2/3-view structure from motion [8,16], rolling shutter camera motion estimation.
- Proposed to estimate motion parameters via a non-convex non-smooth  $\ell^1$  optimization problem, which provably tolerates  $O(\#\text{inliers}^2)$  outliers [9], and performs competitively with state-of-the-art RANSAC algorithms while using an order of magnitude less running time on large-scale SLAM datasets [8,9,16]. Proved the uniqueness of solution via abstract mathematical methods [8].

## TEACHING

---

As a teaching assistant:

ESE 6450 Deep Generative Models	Fall 2025, UPenn
ESE 6450 Deep Generative Models	Fall 2024, UPenn
SI 132 Linear Algebra for Information Science	Spring 2020, ShanghaiTech
SI 232 Subspace Learning	Fall 2019, ShanghaiTech
SI 231 Matrix Analysis	Fall 2018, ShanghaiTech

As a tutor:

AP Calculus, Physics, Computer Science	2017 - 2020, Shanghai American School
Math and Physics of the Chinese College Entrance Exam	2014 - 2017

## TALKS

---

Parsimonious Representations for Analysis and Synthesis	
GRASP SFI ( <b>Winner of the 2025 vote for internal Speakers</b> ), UPenn	2025
Guest lecture invited by the Astralis club, Hamilton High School	2025
Warren & ASSET Center Research Mixer, UPenn	2024
Unsupervised Manifold Linearizing and Clustering	
Conference on Parsimony and Learning (CPAL), Hong Kong University	2024
Vision Lab Retreat, UPenn	2023
Third Workshop on Seeking Low-Dimensionality in Deep Neural Networks (SlowDNN), MBZUAI	2023
Hyperplane Arrangement Descent	
Conference on Parsimony and Learning (CPAL), Hong Kong University	2024
Doubly Stochastic Clustering: Algorithms, Theory, and Applications	
Guest lectures invited by Dr. Chun-Guang Li, Online	2022 & 2023
Mathematical Institute for Data Science (MINDS) Retreat, JHU	2022
International Conference on Machine Learning (ICML), Maryland	2022

## SERVICE

---

Coordinator, PennAI Symposium, UPenn	2025
Coordinator, Conference on the Mathematical Theory of Deep Neural Networks (DeepMath)	2024
Program Committee, Annual AAAI Conference on Artificial Intelligence (AAAI)	2025
Reviewer, International Conference on Learning Representations (ICLR)	2025
Reviewer, European Conference on Computer Vision (ECCV)	2024
Reviewer, International Conference on Acoustics, Speech, and Signal Processing (ICASSP)	2024
Reviewer, International Conference on Machine Learning (ICML)	2022, 2024
Reviewer, Conference on Neural Information Processing Systems (NeurIPS)	2022 - 2024
Reviewer, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)	2023
Reviewer, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)	2022
Reviewer, International Journal of Computer Vision (IJCV)	2021

## REFERENCES

---

<b>René Vidal</b> (PhD advisor) PennAI Co-Chair & Director of IDEAS; PIK & Rachleff University Professor University of Pennsylvania	vidalr@seas.upenn.edu
<b>Yi Ma</b> Professor, Electrical Engineering and Computer Sciences, UC Berkeley Chair of Artificial Intelligence; Director of School of Computing and Data Science, HKU	yima@eecs.berkeley.edu

**Kostas Daniilidis**

Ruth Yalom Stone Professor of Computer and Information Science  
University of Pennsylvania

kostas@cis.upenn.edu

**Chris Callison-Burch**

Professor of Computer and Information Science  
University of Pennsylvania

ccb@upenn.edu

**Jeremias Sulam**

William R. Brody Faculty Scholar; Assistant Professor in Biomedical Engineering  
Johns Hopkins University

jsulam1@jhu.edu

**Manolis C. Tsakiris** (Research advisor during my time at ShanghaiTech)  
Tenured Associate Professor, Academy of Mathematics and Systems Science  
Chinese Academy of Sciences, Beijing

manolis@amss.ac.cn