

San Jose, CA (US)
fangzhoulinark@gmail.com
linkedin.com/in/fangzhou-ark-lin
(774) 823-9051

FANGZHOU LIN

Researcher with a Ph.D. in Computer Science and Data Science, and 5+ years of academic and industrial research in intelligent systems. Develops learning-based perception, predictive world models and planning, 3D reconstruction, and computational imaging from mathematical formulation through reproducible Python/C++/PyTorch implementation and publication. Research appears at CVPR, ICLR, NeurIPS, ICCV, IROS, WACV, and TPAMI, with applied validation in autonomous driving and biomedical imaging. Prepared to build a student-engaged program connecting lectures, laboratories, and capstone projects with research in robotics, signal/image processing, and trustworthy AI.

ACADEMIC FOCUS

Research: Robotics and autonomous systems; perception and planning; world models; 3D reconstruction; signal/image processing; multimodal AI
Teaching Interests: Machine learning; computer vision; robotics; digital image and signal processing; engineering probability and optimization; senior design
Experimental & Computational: Python; C/C++; CUDA; PyTorch; Linux; Docker; simulation; GPU experimentation; reproducible benchmarking

RESEARCH EXPERIENCE

MAY 2026 – PRESENT

Research Scientist, Robert Bosch LLC, Sunnyvale, CA, USA

- Conceived a hyperparameter-free log-cosh objective for latent world-model prediction and planning; across 7 environments and two model frameworks, it never underperformed MSE and improved task success by up to 10 percentage points with geometric regularization.
- Led the project from research question and mathematical formulation through cross-framework implementation, controlled ablations, and failure analysis, producing reproducible guidance for learning-based prediction and planning.

MAY 2025 – AUG 2025

Research Scientist, Robert Bosch LLC, Sunnyvale, CA, USA

- Developed NexusFlow, a lightweight partially supervised framework unifying dense map reconstruction and multi-object tracking under geographic domain shift; achieved state of the art on nuScenes and consistent segmentation, depth, and surface-normal gains on NYUv2 (CVPR 2026).
- Co-developed Pantheon360, a 3D-aware 360° video diffusion framework with an explicit 3D Cache for camera-controllable generation, geometric consistency, and simulation-ready digital twins (CVPR 2026).

AUG 2025 – PRESENT

Research Associate, Texas A&M University, College Station, TX, USA

- Built CV-Arena, a 12K-pair benchmark spanning 16 visual tasks and 21 systems, with Active Elo sampling and a logic-gated VLM judge that integrates human preferences into scalable, interpretable evaluation.
- Developed CV-Agent, a closed-loop plan–edit–verify system that translates visual instructions into structured plans, executes operations, diagnoses discrepancies, and iteratively corrects its outputs.
- Developed a falsification-driven research loop in which adversarial agents generate counterexample environments, expose model failures, and refine candidate solutions; its discovered FNG-CE loss outperformed CE and CE+L2 through ImageNet scale.

NOV 2021 – APR 2025

Research Assistant, Worcester Polytechnic Institute, Worcester, MA, USA

- Developed Deep Loss Convexification, integrating adversarial learning with classical iterative optimization for point-cloud registration and multimodal image alignment (TPAMI).
- Developed GPS, a probabilistic distributional similarity for low-data representation learning that improved few-shot image classification and 3D point-cloud completion across architectures and datasets (ICLR 2025).
- Built a view-free multi-branch Transformer for point-cloud completion, combining cross-attention and self-attention to outperform image-guided multimodal systems on ShapeNet-ViPC.
- Developed plug-in geometric objectives and a gradient-matching loss-search pipeline, delivering robust, efficient state-of-the-art 3D reconstruction across PCN and ShapeNet benchmarks (NeurIPS 2023; ICCV 2023; IROS 2024).
- Led SPADE, a hybrid analytical and data-free ML framework for medical-image denoising, validated across simulation, phantom, ex vivo, and in vivo studies with more than 15 dB SNR improvement in high-noise cases (Photoacoustics 2025).

STUDENT-ENGAGED RESEARCH PROGRAM

- Intelligent Robotics & Control:** Data-efficient world models, multi-sensor perception, and robust planning under distribution shift, organized as progressive projects from simulation to system-level evaluation.
- Computational Imaging & Signal Processing:** Physics-informed and statistical learning for reconstruction, denoising, and sensing, extending prior validation across simulation, phantoms, and biological studies.
- Trustworthy Multimodal AI:** Human-in-the-loop evaluation, verifiable agents, and open benchmarks that support collaborative student teams and applied partnerships.

EDUCATION

APR 2021 – JAN 2024

Doctor of Philosophy in Computer Science and Data Science, Tohoku University, Sendai, JAPAN

Advisors: Prof. Kazunori D. Yamada and Prof. Kengo Kinoshita GPA: 4.0/4.0

SELECTED HONORS

Gold Prize, 3rd AI Mathematical Olympiad Challenge (5th of 3,450 globally), 2026 | Tohoku University Doctoral Fellowship, 2021–2024

PUBLICATIONS

1. **Lin, F.**, Wang, Y., Guo, Y., Huang, Z., Huang, X., Zhang, H., ... & Zhang, Z. (2026). NexusFlow: Unifying Disparate Tasks under Partial Supervision via Invertible Flow Networks. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
2. Chen, T.-H., Chen, Y.-H., Tu, T., Lee, J.-Y., Wu, C.-Y., **Lin, F.**, Zhang, H., Paz, D., Huang, X., Guo, Y., Liu, Y.-L., Wang, Y., & Ren, L. (2026). Pantheon360: Taming Digital Twin Generation via 3D-Aware 360° Video Diffusion. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
3. Gao, X., Wu, M., Yang, S., Yu, J., Taghavi, P., **Lin, F.**, & Tu, Z. (2026). The Pulse of Motion: Measuring Physical Frame Rate from Visual Dynamics. In 1st Workshop on Video World Models: Interaction, Memory, Efficiency. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
4. Li, Z., Chen, X., Jiang, L., Hou, D., Yamada, K., **Lin, F.**, Gao, X., & Tu, Z. (2026). Physics-Aware Video Instance Removal Benchmark. In Workshop on Video Generative Models: Benchmarks and Evaluation, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
5. Wu, M., Yang, S., Xing, S., Mishra, A., Dey, S., Gao, X., Zhao, J., Yu, J., **Lin, F.**, Tu, Z. and Lei, Y., 2026. VGBE 2026 Challenge on Image-to-Video Consistent Generation: Methods and Results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
6. Jiang, L., Hou, D., Wang, Y., Su, Y., **Lin, F.**, et al. (2026). A compact Kolmogorov–Arnold network mixer for long-term time series forecasting. Scientific Reports.
7. **Lin, F.**, Li, P., Xu, L., Chen, W., Ge, Q., Xing, S., ... & Tu, Z. (2026). CV-Arena: An Open Benchmark for Instructional Computer Vision Problem Solving with Human-AI Collaborative Preferences. Under review in The Fortieth Annual Conference on Neural Information Processing Systems (NeurIPS) Evaluations and Datasets Track.
8. **Lin, F.**, Xing, S., Li, P., Yang, S., Ge, Q., Yamada, K., Zhang, Z., Zhang, H., & Tu, Z. (2026). CAPS: Cascaded Adaptive Pairwise Selection for Efficient Parallel Reasoning. Under review in The Fortieth Annual Conference on Neural Information Processing Systems (NeurIPS).
9. Jiang, L., Li, Z., Xing, S., Li, P., Takahashi, T., Hou, D., Tu, Z., Yamada, K., & **Lin, F.** (2026). PathCal: State-Aware Reflection-Marker Calibration for Efficient Reasoning. Under review in The Fortieth Annual Conference on Neural Information Processing Systems (NeurIPS).
10. Li, P., **Lin, F.**, Xing, S., Zheng, X., Hong, X., Sun, J., ... & Ni, C. (2026). BibAgent: An Agentic Framework for Traceable Misinformation Detection in Scientific Literature. Under review in Nature.
11. Li, P., Sun, J., **Lin, F.**, Xing, S., Fu, T., Feng, S., ... & Tu, Z. (2026). Traversal-as-Policy: Log-Distilled Gated Behavior Trees as Externalized, Verifiable Policies for Safe, Robust, and Efficient Agents. arXiv preprint.
12. Li, P., **Lin, F.**, Xing, S., Sun, J., Zhang, D., ... & Tu, Z. (2026). Let the abyss stare back adaptive falsification for autonomous scientific discovery. arXiv preprint.
13. Jiang, L., Wang, Y., Su, Y., Xing, S., Chen, W., Zhang, X., ... **Lin, F.** (2025). KANMixer: Can KAN Serve as a New Modeling Core for Long-term Time Series Forecasting? Under review in Transactions on Machine Learning Research (TMLR).
14. **Lin, F.**, Zhang, Z., Liu, H., Morales, J., Zhang, H., Yamada, K., ... & Saligrama, V. (2025). GPS: A Probabilistic Distributional Similarity with Gumbel Priors for Set-to-Set Matching. In The Thirteenth International Conference on Learning Representations (ICLR).
15. **Lin, F.**, Gao, S., Tang, Y., Ma, X., Murakami, R., ..., Soboyejo, W. O., & Zhang, H. K. (2025). Spectroscopic photoacoustic denoising framework using hybrid analytical and data-free learning method. Photoacoustics, 43, 100729.
16. **Lin, F.**, Gao, S., Tang, Y., Ma, X., Zhang, Z., & Zhang, H. K. (2025). Feasibility study of data-free spectroscopic photoacoustic image denoising. In Medical Imaging 2025: Ultrasonic Imaging and Tomography (Vol. 13412, p. 134120J). SPIE.
17. Kubota, A., Kodate, S., Li, Y., **Lin, F.**, Fukuda, H., Baladram, M. S., & Yamada, K. D. (2025). Learning Random Numbers to Realize Appendable Memory System for Artificial Intelligence to Acquire New Knowledge after Deployment. Interdisciplinary Information Sciences, 31(1), 1–11.
18. **Lin, F.**, Dai, Z., Sanku, R., Hou, S., Yamada, K. D., Zhang, H. K., & Zhang, Z. (2025). A Strong View-Free Baseline Approach for Single-View Image Guided Point Cloud Completion. arXiv preprint arXiv:2506.15747.
19. **Lin, F.**, Hou, S., Liu, H., Gao, S., Yamada, K. D., Zhang, H. K., & Zhang, Z. (2024). Hyperbolic chamfer distance for point cloud completion and beyond. arXiv preprint arXiv:2412.17951.
20. **Lin, F.**, Liu, H., Zhou, H., Hou, S., Yamada, K. D., Fischer, G. S., ... & Zhang, Z. (2024). Loss Distillation via Gradient Matching for Point Cloud Completion with Weighted Chamfer Distance. In 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 511–518).
21. Zhang, Z., Shao, Y., Zhang, Y., **Lin, F.**, Zhang, H., & Rundensteiner, E. (2024). Deep Loss Convexification for Learning Iterative Models. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI).
22. Yue, Y., **Lin, F.**, Mou, G., & Zhang, Z. (2024). Understanding hyperbolic metric learning through hard negative sampling. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV).
23. Yu, X., **Lin, F.**, Yue, Y., & Zhang, Z. (2024). Leveraging Superfluous Information in Contrastive Representation Learning. arXiv preprint arXiv:2408.10292.
24. **Lin, F.**, Yue, Y., Zhang, Z., Hou, S., Yamada, K., Kolachalama, V., & Saligrama, V. (2023). InfoCD: a contrastive chamfer distance loss for point cloud completion. Advances in Neural Information Processing Systems (NeurIPS).
25. **Lin, F.**, Yue, Y., Hou, S., Yu, X., Xu, Y., Yamada, K. D., & Zhang, Z. (2023). Hyperbolic chamfer distance for point cloud completion. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV).
26. Yue, Y., **Lin, F.**, Yamada, K. D., & Zhang, Z. (2023). Hyperbolic contrastive learning. arXiv preprint.
27. **Lin, F.**, Xu, Y., Zhang, Z., Gao, C., & Yamada, K. D. (2022). Cosmos Propagation Network: Deep learning model for point cloud completion. Neurocomputing, 507, 221–234.
28. Yamada, K. D., Baladram, M. S., & **Lin, F.** (2022). Relation is an option for processing context information. Frontiers in Artificial Intelligence, 5, 924688.
29. Gao, C., Cai, G., Jiang, X., Zheng, F., Zhang, J., Gong, Y., **Lin, F.**, Sun, X., & Bai, X. (2022). Conditional feature learning based transformer for text-based person search. IEEE Transactions on Image Processing, 31, 6097–6108.
30. Xu, Y., Arai, S., Liu, D., **Lin, F.**, & Kosuge, K. (2022). FPCC: Fast point cloud clustering-based instance segmentation for industrial bin-picking. Neurocomputing, 494, 255–268.
31. Yamada, K. D., Baladram, M. S., & **Lin, F.** (2022). Progress in research on implementing machine consciousness. Interdisciplinary Information Sciences, 28(1), 95–105.
32. **Lin, F.**, Gao, C., & Yamada, K. D. (2022). An effective convolutional neural network for visualized understanding transboundary air pollution based on Himawari-8 satellite images. IEEE Geoscience and Remote Sensing Letters, 19, 1–5.
33. Yamada, K. D., **Lin, F.**, & Nakamura, T. (2021). Developing a novel recurrent neural network architecture with fewer parameters and good learning performance. Interdisciplinary Information Sciences, 27(1), 25–40.
34. Xu, Y., Arai, S., Liu, D., **Lin, F.**, & Kosuge, K. (2020). FPCC: Fast point cloud clustering for instance segmentation. arXiv preprint arXiv:2012.14618.