

Zexi Jia

✉ zxjia10@gmail.com ☎ (+86) 183-0182-3679 📍 Beijing, China 🎓 Google Scholar 🐙 Github

Research Summary

I work on **trustworthy generative vision**, with broader interests in **multimodal foundation models** and **visual representation learning**. My research develops principled methods for making generative visual systems more **reliable, measurable, and controllable**, including work on **media authenticity and copyright**, **generative model evaluation**, and **large-scale training and optimization**. Across these topics, I seek to translate insights from **geometry, statistics, and visual structure** into methods that advance both scientific understanding and real-world generative systems.

Research Interests

Computer Vision; Generative Models; Trustworthy AI; Multimodal Learning; Biometrics; Event-Camera.

Education

Peking University

Ph.D. in Intelligent Science (Machine Vision)

National Key Laboratory of Visual and Auditory Information Processing

Advisor: Prof. Jufu Feng

Beijing, China

Sep 2020 – Jul 2024

Fudan University

M.Sc. in Computer Science (Computer Vision)

Shanghai Key Laboratory of Data Science

Advisor: Prof. Xiangdong Zhou

Shanghai, China

Sep 2017 – Jul 2020

Fudan University

B.Eng. in Software Engineering

Shanghai Key Laboratory of Data Science

Shanghai, China

Sep 2011 – Jul 2016

Appointments

Tencent WeChat AI

Senior Researcher

Beijing, China

Jul 2024 – Present

Peking University

Ph.D. Researcher

Beijing, China

Sep 2020 – Jul 2024

Microsoft Research Asia

Research Intern

Beijing, China

Jun 2019 – Oct 2019

Research Experience

Deepfake Detection and AIGC Governance

Media Authenticity, Copyright, and Provenance, Beijing, China

Tencent WeChat AI

Jul 2023 – Present

- Studied the authenticity of generated images from physical, statistical, and colorimetric perspectives. Developed generalizable detection methods using specular reflection inconsistency in face forgeries and color-distribution shifts under noise perturbation (ICLR 2026; CVPR 2025; TPAMI 2026).
- Investigated copyright and provenance in AI-generated images, covering content, style, composition, and character IP. Developed methods for multi-aspect infringement detection, artistic-style analysis, structural plagiarism detection, and training-free watermarking for diffusion models (ICCV 2025; ICASSP 2025; CVPR 2026).
- Built research prototypes and online systems for large-scale AIGC safety and copyright risk assessment, with emphasis on deployability, efficiency, and robustness.

Perceptual Metrics and Benchmark Design

Generative Model Evaluation, Beijing, China

Tencent WeChat AI

Dec 2024 – Present

- Developed perceptually aligned evaluation methods for generated images, including **ArtFRD**, an aesthetic metric grounded in Fisher–Rao geometry over image representation distributions (ACM MM 2025, Oral).
- Proposed **CHD/CMMS**, metrics based on one-dimensional distributions of discrete visual tokens, achieving strong correlation with human perceptual preferences in generative model evaluation (CVPR 2026, Highlight).
- Built evaluation protocols for realism, aesthetic quality, and color fidelity, including benchmark design for calibrating saturation bias in photorealistic image generation (CVPR 2026, Highlight).

Training and Optimization of Multimodal Systems

Large-Scale Generative Models, Beijing, China

Tencent WeChat AI

Jul 2024 – Present

- Led a **7-person team** on the **CrownFish** image generation project (10B/16B), including architecture design, training strategy, and large-scale experimentation on a **512×H100** GPU cluster.
- Built data, evaluation, and regression pipelines for **WeClip (8B)** and **WeMM (32B)**, and led work on training stability, cross-modal alignment, and reasoning-oriented optimization.
- Studied guidance-induced artifacts in diffusion models and proposed **Manifold-Optimal Guidance (MOG)**, a sampling correction framework derived from a Riemannian optimal control perspective (ICML 2026, under review).
- Supported deployment of generative and multimodal systems to large-scale WeChat products, including sticker generation, retrieval, and e-commerce applications.

Fingerprint Analysis and Large-Scale Matching

Biometric Recognition, Beijing, China

Peking University

Sep 2020 – Jul 2023

- Served as lead algorithm developer for a next-generation national fingerprint identification system supporting billion-scale daily identity matching.
- Developed algorithms for fingerprint segmentation, restoration, enhancement, and minutiae matching under challenging acquisition conditions, achieving a **65%** accuracy improvement over the previous system in official Ministry of Public Security evaluations.
- Published research in *IEEE TIFS*, *IEEE TIP*, *IEEE SPL*, and *IJCB*, covering incomplete fingerprint identification, latent fingerprint segmentation, and minutiae restoration.

Attention Mechanisms and Semantic Segmentation

Event-Based Vision, Beijing, China

Peking University

Oct 2021 – Jun 2023

- Studied event-camera perception for autonomous driving, with a focus on **semantic segmentation**, **pre-training**, and efficient visual representation learning under sparse and asynchronous sensing.
- Proposed the **Posterior Attention** mechanism and an event-camera pre-training framework, advancing state-of-the-art performance on DDD17 and MVSEC benchmarks.
- Completed the full research loop from problem formulation to deployment, with resulting algorithms deployed on HiSilicon chips and recognized across multiple teams.

Professional Service

- **Reviewer:** CVPR, ICCV, ECCV, ICLR, ICML, NeurIPS, AAAI, ACM MM; IEEE TPAMI, TIP, TIFS.
- **Mentoring:** Mentored **15+** graduate students and research interns on generative models, AIGC safety, and multimodal learning.

Selected Honors & Awards

- Tencent Qingyun Award (Top Young Talent Selection), 2024
- Huawei Top Minds Award, 2023
- Highest-Level Award, Ministry of Public Security, P.R. China, 2023
- Featured in *People's Daily* and *Guangming Daily* for fingerprint recognition breakthroughs, 2023

Publications

* indicates corresponding author.

Selected Publications.....

- **Zexi Jia**, Zhiqiang Yuan, Xiaoyue Duan, Jinchao Zhang*, Jie Zhou, and Anil K. Jain. “CoDA: Color Distribution Probing for Efficient and Generalizable AI-Generated Image Detection”. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, Early Access, **2026**.
- **Zexi Jia**, Pengcheng Luo, Yijia Zhong, Jinchao Zhang, and Jie Zhou. “Evaluating Generative Models via One-Dimensional Code Distributions”. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, **2026**. **(Highlight)**
- **Zexi Jia**, Chuanwei Huang, Yeshuang Zhu, Hongyan Fei, Ying Deng, Zhiqiang Yuan, Jiapei Zhang, Jinchao Zhang*, and Jie Zhou. “A Visual Leap in CLIP Compositionality Reasoning through Generation of Counterfactual Sets”. *IEEE/CVF International Conference on Computer Vision (ICCV)*, **2025**.
- **Zexi Jia**, Chuanwei Huang, Yeshuang Zhu, Hongyan Fei, Xiaoyue Duan, Zhiqiang Yuan, Ying Deng, Jiapei Zhang, Jinchao Zhang*, and Jie Zhou. “Secret Lies in Color: Enhancing AI-Generated Images Detection with Color Distribution Analysis”. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, **2025**.
- **Zexi Jia**, K. You, W. He, Y. Tian, Y. Feng, Y. Wang, X. Jia, Y. Lou, J. Zhang, and G. Li*. “Event-based Semantic Segmentation with Posterior Attention”. *IEEE Transactions on Image Processing (TIP)*, **2023**.

Other First-Author and Corresponding-Author Publications.....

- Zhengyao Fang, **Zexi Jia***, Yijia Zhong, Pengcheng Luo, Jinchao Zhang, Guangming Lu, Jun Yu, and Wenjie Pei. “Too Vivid to Be Real? Benchmarking and Calibrating Generative Color Fidelity.” *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, **2026**. **(Highlight)**
- Pengcheng Luo, **Zexi Jia***, Yijia Zhong, Jinchao Zhang, and Jie Zhou. “GROW: Watermark Generation with Progressive Guidance for Diffusion Models.” *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, **2026**.
- Hongyan Fei, **Zexi Jia***, Chuanwei Huang, Jinchao Zhang, and Jie Zhou. “Exploring Specular Reflection Inconsistency for Generalizable Face Forgery Detection.” *International Conference on Learning Representations (ICLR)*, **2026**.
- **Zexi Jia**, Jinchao Zhang, and Jie Zhou. “StyleDecoupler: Generalizable Artistic Style Disentangle.” *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, **2026**.
- **Zexi Jia**, Chuanwei Huang, Yeshuang Zhu, Hongyan Fei, Zhiqiang Yuan, Ying Deng, Jiapei Zhang, Jinchao Zhang*, and Jie Zhou. “From Imitation to Innovation: The Emergence of AI’s Unique Artistic Styles and the Challenge of Copyright Protection.” *IEEE/CVF International Conference on Computer Vision (ICCV)*, **2025**.
- Chuanwei Huang, **Zexi Jia***, Hongyan Fei, Yeshuang Zhu, Zhiqiang Yuan, Ying Deng, Jiapei Zhang, Xiaoyue Duan, Jinchao Zhang, and Jie Zhou. “MCID: Multi-aspect Copyright Infringement Detection for Generated Images.” *IEEE/CVF International Conference on Computer Vision (ICCV)*, **2025**.
- Chuanwei Huang, **Zexi Jia***, Yeshuang Zhu, Hongyan Fei, Jinchao Zhang, and Jie Zhou. “ArtFRD: A Fisher–Rao Mixture Metric for Generative Model Aesthetic Evaluation.” *ACM International Conference on Multimedia (ACM MM)*, **2025**. **(Oral)**
- Chuanwei Huang, **Zexi Jia***, Hongyan Fei, Yeshuang Zhu, Zhiqiang Yuan, Ying Deng, Jiapei Zhang, Fei Jiao, Jinchao Zhang, and Jie Zhou. “Semantic to Structure: Learning Structural Representations for Infringement Detection.” *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, **2025**. **(Oral)**
- **Zexi Jia*** and Chuanwei Huang. “Automated Framework for Extracting and Restoring Minutiae from Low-Quality Fingerprints.” *IEEE Signal Processing Letters (SPL)*, **2025**.

- **Zexi Jia**, Chuanwei Huang, Zheng Wang, Hongyan Fei, Song Wu, and Jufu Feng*. “Finger Recovery Transformer: Towards Better Incomplete Fingerprint Identification.” *IEEE Transactions on Information Forensics and Security (TIFS)*, **2024**.
- **Zexi Jia**, Chuanwei Huang, Zheng Wang, and Jufu Feng*. “FingerSTR: Weak Supervised Transformer for Latent Fingerprint Segmentation.” *IEEE International Joint Conference on Biometrics (IJCB)*, **2023**.

Co-authored Publications.....

- Hanbo Bi, Zhiqiang Yuan, **Zexi Jia**, Jiawei Zhang, Chongyang Li, Peixiang Luo, Ying Deng, Xiaoyue Duan, and Jinchao Zhang*. “F²RVLM: Boosting Fine-grained Fragment Retrieval for Multi-Modal Long-form Dialogue with Vision Language Model.” *AAAI Conference on Artificial Intelligence (AAAI)*, **2026**.
- Chongyang Li, Zhiqiang Yuan, Hanbo Bi, **Zexi Jia**, and Jinchao Zhang*. “Less Redundancy: Boosting Practicality of Vision Language Model in Walking Assistants.” *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, **2026**.
- Zhiqiang Yuan, Ting Zhang, Yeshuang Zhu, Ying Deng, Jiawei Zhang, **Zexi Jia**, Xiaoyue Duan, Jinchao Zhang, and Jie Zhou*. “WalkVLM: Aid Visually Impaired People Walking by Vision Language Model.” *IEEE/CVF International Conference on Computer Vision (ICCV)*, **2025**.
- Hongyan Fei, Chuanwei Huang, Song Wu, Zheng Wang, **Zexi Jia**, and Jufu Feng*. “Fingerprint Presentation Attack Detection by Region Decomposition.” *IEEE Transactions on Information Forensics and Security (TIFS)*, **2024**.
- Chuanwei Huang, Hongyan Fei, **Zexi Jia**, Song Wu, and Jufu Feng*. “Fingerprint Presentation Attack Detection with Supervised Contrastive Learning.” *IEEE International Joint Conference on Biometrics (IJCB)*, **2023**.
- Song Wu, **Zexi Jia**, and Jufu Feng*. “Minutiae-awarely Learning Fingerprint Representation for Fingerprint Indexing.” *IEEE International Joint Conference on Biometrics (IJCB)*, **2022**. (Oral)
- Zheng Wang, **Zexi Jia**, and Jufu Feng*. “A Semantic Segmentation Method for Overlapping Fingerprints.” *IEEE Signal Processing Letters (SPL)*, **2022**.

Manuscripts Under Review.....

- **Zexi Jia**, Pengcheng Luo, Zhengyao Fang, Jinchao Zhang*, and Jie Zhou. “Manifold-Optimal Guidance: A Unified Riemannian Control View of Diffusion Guidance.” *Submitted to the International Conference on Machine Learning (ICML, pre-rebuttal: 5544)*, **2026**.
- **Zexi Jia**, Jinchao Zhang, and Jie Zhou. “From Clusters to Generators: Cycle-Consistent Unsupervised Concept Learning.” *Submitted to the European Conference on Computer Vision (ECCV)*, **2026**.

Patents

- **16 granted Chinese invention patents** as first inventor, covering AIGC safety and copyright governance (forgery detection, sampling optimization, automated quality evaluation) and large-scale biometric recognition (fingerprint segmentation, restoration, high-performance retrieval, and liveness detection).
- **1 U.S. patent application pending** as first inventor, related to AIGC safety and copyright governance.