Zhipeng Wei

Postdoc at International Computer Science Institute
University of California, Berkeley
Berkeley, California, USA

zwei@icsi.berkeley.edu
zhipeng-wei.github.io
Google Scholar

Work Experience

International Computer Science Institute, UC Berkeley *Berkeley, USA*
Postdoc Research Fellow *Nov.2024 - present*
Advisor: Prof. N. Benjamin Erichson

School of Computer Science, Fudan University *Shanghai, China*
Research Assistant *Jul.2020 – Aug.2021*
Advisor: Prof. Chen Jingjing and Jiang Yu-Gang

School of Computing, National University of Singapore *Singapore*
Visiting Research Scholar *Oct.2018 – Jun.2019*
Advisor: Prof. Chua Tat Seng

Education

Fudan University *Shanghai, China*
Ph.D. in Computer Science *Sep.2021- Jun.2024*

- Thesis: Adversarial Robustness in Deep Visual Models
- Supervisor: Prof. Chen Jingjing and Jiang Yu-Gang

Jilin University *Changchun, China*
M.S. in Computer Science *Sep.2017 – Jun.2020*

- Supervisor: Prof. Zhou Fengfeng

Jilin University *Changchun, China*
B.S. in Biological Science *Sep.2013 - Jun.2017*

Research Interests

My research interests include advancing the robustness and trustworthiness of generative foundation models, with a focus on ensuring reliable, ethical, and resilient model outputs.

TEACHIING EXPERIENCE

Fudan University *Shanghai, China*

Teaching Assistant *2021 - 2023*

- Taught over 50 undergraduate students in the computer vision course, delivering 135-minute sessions twice per semester, which included OpenCV and PyTorch tutorials.

Fudan University *Shanghai, China*

Mentoring *2021 - 2024*

- Assisted in mentoring a team focused on adversarial robustness, with published papers in AAAI, ACM MM, and ICME.

SELECTED AWARDS AND HONORS

Shanghai Outstanding Graduates	*Jun.2024*
Outstanding Master's Thesis of Jilin Province	*Dec.2022*
China National Scholarship of Fudan University	*Oct.2022*
Outstanding Graduates of Jilin University	*Jun.2020*
Outstanding Master's Thesis of Jilin University	*Jun.2020*

ACADEMIC SERVICE

Reviewer for IEEE TIP; TNNLS; TCSVT; TDSC; ACM MM 2022; AAAI 2023, 2024, 2025; CVPR 2023, 2024; IJCAI 2023; ICCV 2023; NIPS2023, 2024; ICLR 2024, 2025; ECCV 2024.

PUBLICATIONS

1. Wenzhuo Xu, Kai Chen, Ziyi Gao, **Zhipeng Wei**, Jingjing Chen, and Yu-Gang Jiang, "Highly transferable diffusion-based unrestricted adversarial attack on pre-trained vision-language models", In: Proceedings of the 32nd ACM International Conference on Multimedia. 2024
2. Ziyi Gao, Kai Chen, **Zhipeng Wei**, Tingshu Mou, Jingjing Chen, Zhiyu Tan, Hao Li, and Yu-Gang Jiang, "ReToMe-VA: Recursive Token Merging for Video Diffusion-based Unrestricted Adversarial Attack", In: Proceedings of the 32nd ACM International Conference on Multimedia. 2024
3. Chao Gong, Kai Chen, **Zhipeng Wei**, Jingjing Chen, and Yu-Gang Jiang, "Reliable and efficient concept erasure of text-to-image diffusion models", In: Proceedings of the 32nd ACM International Conference on Multimedia. 2024

4. Yan Jiang, Guisheng Yin, Ye Yuan, Jingjing Chen, and **Zhipeng Wei**, "Cross-Point Adversarial Attack Based on Feature Neighborhood Disruption Against Segment Anything Model". In: 2024 IEEE International Conference on Multimedia and Expo (ICME). IEEE. 2024, pp. 1-6
5. **Zhipeng Wei**, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang, "Adaptive Cross-Modal Transferable Adversarial Attacks from Images to Videos". In: IEEE Transactions on Pattern Analysis and Machine Intelligence. doi: 10.1109/TPAMI.2023.3347835.
6. **Zhipeng Wei**, Jingjing Chen, Micah Goldblum, Zuxuan Wu, Tom Goldstein, Yu-Gang Jiang and Larry S. Davis, "Towards Transferable Adversarial Attacks on Image and Video Transformers". In: IEEE Transactions on Image Processing. vol. 32, pp. 6346-6358, 2023.
7. Kai Chen, **Zhipeng Wei**, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. "GCMA: Generative Cross-Modal Transferable Adversarial Attacks from Images to Videos." In: Proceedings of the 31th ACM International Conference on Multimedia. 2023
8. Jingjing Chen*, Linhai Zhuo*, **Zhipeng Wei**, Hao Zhang, Huazhu Fu, and Yu-Gang Jiang. "Knowledge driven weights estimation for large-scale few-shot image recognition." In: Pattern Recognition 142 (2023), p. 109668
9. Yiqiang Lv, Jingjing Chen, **Zhipeng Wei**, Kai Chen, Zuxuan Wu, and Yu-Gang Jiang. "Downstream Task-agnostic Transferable Attacks on Language-Image Pre-training Models." In: *2023 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE. 2023.
10. **Zhipeng Wei**, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. "Enhancing the Self-Universality for Transferable Targeted Attacks." In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023, pp. 12281–12290
11. Kai Chen, **Zhipeng Wei**, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. "Attacking video recognition models with bullet-screen comments." In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. 1. 2022, pp. 312–320
12. **Zhipeng Wei**, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. "Cross-modal transferable adversarial attacks from images to videos." In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022, pp. 15064–15073
13. **Zhipeng Wei**, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. "Boosting the transferability of video adversarial examples via temporal translation." In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. 3. 2022, pp. 2659–2667
14. **Zhipeng Wei**, Jingjing Chen, Micah Goldblum, Zuxuan Wu, Tom Goldstein, and Yu-Gang Jiang. "Towards transferable adversarial attacks on vision transformers." In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. 3. 2022, pp. 2668–2676
15. **Zhipeng Wei**, Jingjing Chen, Hao Zhang, Linxi Jiang, and Yu-Gang Jiang. "Adaptive Temporal Grouping for Black-box Adversarial Attacks on Videos." In: *Proceedings of the 2022 International Conference on Multimedia Retrieval*. 2022, pp. 587–593

16. **Zhipeng Wei**, Jingjing Chen, Xingxing Wei, et al. “Heuristic black-box adversarial attacks on video recognition models.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. 07. 2020, pp. 12338–12345
17. Jingjing Chen, Liangming Pan, **Zhipeng Wei**, Xiang Wang, Chong-Wah Ngo, and Tat-Seng Chua. “Zero-shot ingredient recognition by multi-relational graph convolutional network.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. 07. 2020, pp. 10542–10550
18. **Zhipeng Wei**, Jingjing Chen, Zhaoyan Ming, Chong-Wah Ngo, Tat-Seng Chua, and Fengfeng Zhou. “DietLens-eout: Large scale restaurant food photo recognition.” In: *Proceedings of the 2019 on International Conference on Multimedia Retrieval*. 2019, pp. 399–403

INVITED TALKS

Spotlight, Vision And Learning Seminar (VALSE)	*Aug.2022*
Speaker, AI TIME PHD-CVPR	*Aug.2022*
Speaker, AI TIME PHD-AAAI	*May.2022*