

JUNLIN HAN

Webpage · Email · Google Scholar

Last Updated: April 10, 2026

EDUCATION

University of Oxford PhD (DPhil) in Computer Vision Supervised by Prof. Philip Torr	01.2024 – ongoing Oxford, UK
Australian National University Bachelor of Information Technology (Honours) GPA: 3.95/4.00	02.2019 – 07.2023 Canberra, Australia First-Class Honours

EXPERIENCE

Meta <i>PhD Student Researcher</i> <ul style="list-style-type: none">FAIR (12.2025 – ongoing) <i>Host: Dr. Mike Lewis; Topics: Unified pre-training</i>GenAI Llama Vision (08.2024 – 12.2025) <i>Host: Dr. Filippos Kokkinos; Topics: MLLM post-training and pre-training</i>GenAI RepliGen (10.2023 – 07.2024) <i>Host: Dr. Filippos Kokkinos; Topics: 3D Generation, video generation</i>	10.2023 – ongoing London, UK
Cybever Lead Research Scientist Topic: Generative models for 3D world modeling	05.2023 – 09.2023 Remote
AIML, University of Adelaide Visiting Researcher Advisor: Prof. Ian Reid Topics: Data-centric AI, memorability	12.2021 – 05.2023 Adelaide, Australia
Data61-CSIRO & Australian National University Research Student Advisors: Dr. Lars Petersson and Prof. Hongdong Li Topics: Generative models, self-supervised learning	08.2020 – 05.2023 Canberra, Australia

SELECTED PUBLICATIONS

See my [webpage](#) for full publications, codes, datasets, and project pages.

- Learning to See Before Seeing: Demystifying LLM Visual Priors from Language Pre-training**
Junlin Han, Shengbang Tong, David Fan, Yufan Ren, Koustuv Sinha, Philip Torr, Filippos Kokkinos
International Conference on Learning Representations (ICLR oral), 2026
- Flex3D: Feed-Forward 3D Generation with Flexible Reconstruction Model and Input View Curation**
Junlin Han, Jianyuan Wang, Andrea Vedaldi, Philip Torr, Filippos Kokkinos
International Conference on Machine Learning (ICML), 2025
- VFusion3D: Learning Scalable 3D Generative Models from Video Diffusion Models**
Junlin Han*, Filippos Kokkinos*, Philip Torr
European Conference on Computer Vision (ECCV), 2024
- What Images are More Memorable to Machines?**
Junlin Han, Huangying Zhan, Jie Hong, Pengfei Fang, Hongdong Li, Lars Petersson, Ian Reid
Preprint
- You Only Cut Once: Boosting Data Augmentation with a Single Cut**
Junlin Han, Pengfei Fang, Weihao Li, Jie Hong, Ali Armin, Ian Reid, Lars Petersson, Hongdong Li
International Conference on Machine Learning (ICML), 2022

1. Blind Image Decomposition

Junlin Han, Weihao Li, Pengfei Fang, Chunyi Sun, Jie Hong, Ali Armin, Lars Petersson, Hongdong Li
European Conference on Computer Vision (ECCV), 2022

CONTRIBUTED PROJECTS

• The Llama4 herd: The beginning of a new era of natively multimodal AI innovation

- Lead the effort for spatial intelligence in Llama4 visual post-training.
- Created a large-scale (over 100K) internal 3D multi-image understanding dataset covering multiple key 3D understanding abilities like depth estimation, correspondence matching, and multi-view reasoning.
- Designed a pipeline to label multiview raw data using 3D vision methods for precise annotation, which included an automated validation process.
- Integrated the dataset into Llama4 training workflows with rejection sampling, boosting multi-view understanding performance by over 30% on key benchmarks.

ACADEMIC SERVICES

- **Workshop Organization:** Foundation Models in the Wild (ICML 2024, ICLR 2025)
- **Conference Review:** CVPR (2022-2026), ICCV (2023, 2025), ECCV (2022, 2024), ICML (2022, 2024, 2025), NeurIPS (2022, 2023, 2025), ICLR (2023-2026), COLM (2025), SIGGRAPH Asia (2024, 2025)
- **Journal Review:** TPAMI, IJCV, TIP

AWARDS

Meta PhD Scholarship <i>for PhD research at University of Oxford</i>	2023
Chancellor's Letter of Commendation <i>for best academic performance at Australian National University</i>	2022
Top Reviewer <i>NeurIPS (Conference on Neural Information Processing Systems)</i>	2022
Second Best Presentation Award <i>AIM (Active Integrated Matter) Conference</i>	2021
Top-up Scholarship <i>for research work at Data61-CSIRO</i>	2021
Undergraduate Vacation Scholarship <i>for summer research at Data61-CSIRO</i>	2020

FULL PUBLICATIONS (* DENOTES EQUAL CONTRIBUTION)

30. **Small Vision-Language Models are Smart Compressors for Long Video Understanding**
Junjie Fei, Jun Chen, Zechun Liu, Yunyang Xiong, Chong Zhou, Wei Wen, **Junlin Han**, Mingchen Zhuge, Saksham Suri, Qi Qian, Shuming Liu, Lemeng Wu, Raghuraman Krishnamoorthi, Vikas Chandra, Mohamed Elhoseiny, Chenchen Zhu
Tech report, 2026
29. **Beyond Language Modeling: An Exploration of Multimodal Pretraining**
Shengbang Tong, David Fan, John Nguyen, Ellis Brown, Gaoyue Zhou, Shengyi Qian, Boyang Zheng, Théophile Vallaeys, **Junlin Han**, Rob Fergus, Naila Murray, Marjan Ghazvininejad, Mike Lewis, Nicolas Ballas, Amir Bar, Michael Rabbat, Jakob Verbeek, Luke Zettlemoyer, Koustuv Sinha, Yann LeCun, Saining Xie
Text report, 2026
28. **EgoEdit: Dataset, Real-Time Streaming Model, and Benchmark for Egocentric Video Editing**
Runjia Li, Moayed Haji-Ali, Ashkan Mirzaei, Chaoyang Wang, Arpit Sahni, Ivan Skorokhodov, Aliaksandr Siarohin, Tomas Jakab, **Junlin Han**, Sergey Tulyakov, Philip Torr, Willi Menapace
Conference on Computer Vision and Pattern Recognition (CVPR), 2026
27. **From Pixels to Feelings: Aligning MLLMs with Human Cognitive Perception of Images**
Yiming Chen*, **Junlin Han***, Tianyi Bai, Shengbang Tong, Filippos Kokkinos, Philip Torr
Preprint

26. **Scaling Sequence-to-Sequence Generative Neural Rendering**
Shikun Liu, Kam Woh Ng, Wonbong Jang, Jiadong Guo, **Junlin Han**, Haozhe Liu, Yiannis Douratsos, Juan C Pérez, Zijian Zhou, Chi Phung, Tao Xiang, Juan-Manuel Pérez-Rúa
International Conference on Learning Representations (ICLR), 2026
25. **Learning to See Before Seeing: Demystifying LLM Visual Priors from Language Pre-training**
Junlin Han, Shengbang Tong, David Fan, Yufan Ren, Koustuv Sinha, Philip Torr, Filippos Kokkinos
International Conference on Learning Representations (ICLR oral). 2026
24. **Unsupervised Decomposition of 3D Shapes into Expressive and Editable Extruded Profile Primitives**
Chunyi Sun, **Junlin Han**, Runjia Li, Wejian Deng, Dylan Campbell, Stephen Gould
ACM SIGGRAPH, 2025
23. **Flex3D: Feed-Forward 3D Generation With Flexible Reconstruction Model And Input View Curation**
Junlin Han, Jianyuan Wang, Andrea Vedaldi, Philip Torr, Filippos Kokkinos
International Conference on Machine Learning (ICML), 2025
22. **Hallucination at a Glance: Controlled Visual Edits and Fine-Grained Multimodal Learning**
Tianyi Bai, Yuxuan Fan, Jiantao Qiu, Fupeng Sun, Jiayi Song, **Junlin Han**, Zichen Liu, Conghui He, Wentao Zhang, Binhang Yuan
NeurIPS, 2025
21. **AuralSAM2: Enabling SAM2 Hear Through Pyramid Audio-Visual Feature Prompting**
Yuyuan Liu, Yuanhong Chen, Chong Wang, **Junlin Han**, Junde Wu, Can Peng, Jingkun Chen, Yu Tian, Gustavo Carneiro
Conference on Computer Vision and Pattern Recognition Findings (CVPR Findings), 2026
20. **Reliable and Responsible Foundation Models**
Xinyu Yang*, **Junlin Han***, Rishi Bommasani, Jinqi Luo, Wenjie Qu, Wangchunshu Zhou, Adel Bibi, Xiyao Wang, Jaehong Yoon, Elias Stengel-Eskin, Shengbang Tong, Lingfeng Shen, Rafael Rafailov, Runjia Li, Zhaoyang Wang, Yiyang Zhou, Chenhang Cui, Yu Wang, Wenhao Zheng, Huichi Zhou, Jindong Gu, Zhaorun Chen, Peng Xia, Tony Lee, Thomas P Zollo, Vikash Sehwal, Jixuan Leng, Jiuhai Chen, Yuxin Wen, Huan Zhang, Zhun Deng, Linjun Zhang, Pavel Izmailov, Pang Wei Koh, Yulia Tsvetkov, Andrew Gordon Wilson, Jiaheng Zhang, James Zou, Cihang Xie, Hao Wang, Philip Torr, Julian McAuley, David Alvarez-Melis, Florian Tramèr, Kaidi Xu, Suman Jana, Chris Callison-Burch, Rene Vidal, Filippos Kokkinos, Mohit Bansal, Beidi Chen, Huaxiu Yao
Transactions on Machine Learning Research (TMLR), 2025
19. **VGRP-Bench: Visual Grid Reasoning Puzzle Benchmark for Large Vision-Language Models**
Yufan Ren, Konstantinos Tertikas, Shalini Maiti, **Junlin Han**, Tong Zhang, Sabine Süssstrunk, Filippos Kokkinos
Preprint
18. **Generalized Few-shot 3D Point Cloud Segmentation with Vision-Language Model**
Zhaochong An, Guolei Sun, Yun Liu, Runjia Li, **Junlin Han**, Ender Konukoglu, Serge Belongie
Conference on Computer Vision and Pattern Recognition (CVPR), 2025
17. **DreamBeast: Distilling 3D Fantastical Animals with Part-Aware Knowledge Transfer**
Runjia Li, **Junlin Han**, Luke Melas-Kyriazi, Chunyi Sun, Zhaochong An, Zhongrui Gui, Shuyang Sun, Philip Torr, Tomas Jakab
International Conference on 3D Vision (3DV), 2025
16. **3D-GPT: Procedural 3D Modeling with Large Language Models**
Chunyi Sun*, **Junlin Han***, Wejian Deng, Xinlong Wang, Zishan Qin, Stephen Gould
International Conference on 3D Vision (3DV), 2025
15. **Semantic Score Distillation Sampling for Compositional Text-to-3D Generation**
Ling Yang, Zixiang Zhang, **Junlin Han**, Bohan Zeng, Runjia Li, Philip Torr, Wentao Zhang
Preprint
14. **VFusion3D: Learning Scalable 3D Generative Models from Video Diffusion Models**
Junlin Han*, Filippos Kokkinos*, Philip Torr
European Conference on Computer Vision (ECCV), 2024
13. **Strong and Controllable Blind Image Decomposition**
Zeyu Zhang*, **Junlin Han***, Chenhui Gou*, Hongdong Li, Liang Zheng
Preprint

12. **How Many Unicorns Are in This Image? A Safety Evaluation Benchmark for Vision LLMs**
Haoqin Tu*, Chenhang Cui*, Zijun Wang*, Yiyang Zhou, Bingchen Zhao, **Junlin Han**, Wangchunshu Zhou, Huaxiu Yao, Cihang Xie
European Conference on Computer Vision (ECCV), 2024
11. **NeRFEditor: Differentiable Style Decomposition for 3D Scene Editing**
Chunyi Sun, Yanbin Liu, **Junlin Han**, Stephen Gould
Winter Conference on Applications of Computer Vision (WACV), 2024
10. **Hyperbolic Audio-visual Zero-shot Learning**
Jie Hong, Zeeshan Hayder, **Junlin Han**, Pengfei Fang, Mehrtash Harandi, Lars Petersson
International Conference on Computer Vision (ICCV), 2023
9. **Curved Geometric Networks for Visual Anomaly Recognition**
Jie Hong, Pengfei Fang, Weihao Li, **Junlin Han**, Lars Petersson, Mehrtash Harandi
IEEE Transactions on Neural Networks and Learning Systems (TNNLS), 2023
8. **What Images are More Memorable to Machines?**
Junlin Han, Huangying Zhan, Jie Hong, Pengfei Fang, Hongdong Li, Lars Petersson, Ian Reid
Preprint
7. **GOSS: Towards Generalized Open-set Semantic Segmentation**
Jie Hong, Weihao Li, **Junlin Han**, Jiyang Zheng, Pengfei Fang, Mehrtash Harandi, Lars Petersson
The Visual Computer, 2023
6. **CropMix: Sampling a Rich Input Distribution via Multi-Scale Cropping**
Junlin Han, Lars Petersson, Hongdong Li, Ian Reid
Preprint
5. **You Only Cut Once: Boosting Data Augmentation with a Single Cut**
Junlin Han, Pengfei Fang, Weihao Li, Jie Hong, Ali Armin, Ian Reid, Lars Petersson, Hongdong Li
International Conference on Machine Learning (ICML), 2022
4. **Blind Image Decomposition**
Junlin Han, Weihao Li, Pengfei Fang, Chunyi Sun, Jie Hong, Ali Armin, Lars Petersson, Hongdong Li
European Conference on Computer Vision (ECCV), 2022
3. **Underwater Image Restoration via Contrastive Learning and a Real-world Dataset**
Junlin Han, Mehrdad Shoeiby, Tim Malthus, Elizabeth Botha, Janet Anstee, Saeed Anwar, Ran Wei, Ali Armin, Hongdong Li, Lars Petersson
Remote Sensing (Best paper award), 2022
2. **Single Underwater Image Restoration by Contrastive Learning**
Junlin Han, Mehrdad Shoeiby, Tim Malthus, Elizabeth Botha, Janet Anstee, Saeed Anwar, Ran Wei, Lars Petersson, Ali Armin
IEEE International Geoscience and Remote Sensing Symposium (oral), 2021
1. **Dual Contrastive Learning for Unsupervised Image-to-Image Translation**
Junlin Han, Mehrdad Shoeiby, Lars Petersson, Ali Armin
CVPR NTIRE Workshops (oral), 2021