

Chenhongyi Yang

✉ [Mail](#) | 🏠 [Homepage](#) | 🐙 [Github](#) | 🔗 [LinkedIn](#) | 📄 [Google Scholar](#)

Education

University of Edinburgh

Mar. 2021 – July 2024

PhD in Engineering

Boston University

Sep. 2018 – May 2020

MSc in Computer Science

University of Science and Technology of China

Sep. 2014 – June 2018

BEng in Computer Science and Technology

Employment

Meta | *Research Scientist Intern*

June 2023 – Mar. 2024

- Egocentric 3D Human Pose Estimation. Developed the first transformer for egocentric pose estimation. It achieved a 42% less error compared to Meta's production model on the internal benchmark. On the public UnrealEgo dataset, it outperformed the previous art by 45% fewer errors with only 13.1% FLOPs. ([EgoPoseFormer](#) - ECCV 2024)

TuSimple | *Research Intern*

May 2020 – Feb 2021

- Efficient Small Object Detection. Introduced a sparse computing mechanism designed to enhance the speed of small object detection using high-resolution feature maps. The mechanism, when tested on the public COCO dataset, allowed the detection model to operate $3.2\times$ faster while maintaining accuracy with minimal loss. ([QueryDet](#) - CVPR 2022)

Research Projects

BayesWatch Group, University of Edinburgh | *Doctoral Researcher*

Mar. 2021 – Present

- General Visual Recognition. (1) Developed a high-resolution, non-hierarchical vision transformer for general visual recognition. The new architecture achieved much better performance than previous models across various downstream tasks. ([GPViT](#) - ICLR 2023); (2) Introduced a simple Mamba architecture as an image encoder. With linear complexity relative to the input size, it outperformed the vision transformer on multiple downstream tasks. ([PlainMamba](#) - BMVC 2024)
- Efficient Training for Visual Recognition Models. (1) Developed a model-independent active learning algorithm for object detection, surpassing the performance of all prior methods. ([PPAL](#) - CVPR 2024); (2) Designed a state-of-the-art and highly generalizable knowledge distillation (KD) approach for object detection, utilizing the teacher model's performance to guide the computation of KD loss. ([PGD](#) - ECCV 2022); (3) Introduced an unsupervised pre-training paradigm for instance-level tasks, employing contrastive loss computed with regional visual features. ([CCOP](#) - Preprint)
- Robust Transformer for BEV View-transformation. Designed a transformer to convert multi-view image features into a unified BEV representation. The transformer processes compressed visual features as inputs and is composed of a single standard cross-attention layer, making it robust, efficient, and deployment-friendly. ([WidthFormer](#) - IROS 2024)

IVC Lab, Boston University | *Research Assistant*

Jan. 2019 – May 2020

- Detecting Heavily Occluded Objects. Introduced a novel non-maximum suppression (NMS) algorithm that significantly enhanced the detection recall for heavily occluded objects. Benchmarking on the KITTI and CityPersons datasets confirmed the algorithm's efficacy. ([SG-NMS](#) - ECCV 2020)

Skills

- **Computer Vision**: Foundation Models, Vision Language Models, Diffusion Models, ConvNets, Transformers, Mamba, Image Classification, Object Detection, Image Segmentation, 3D Object Recognition, 3D Human Pose Estimation, Egocentric Vision, Knowledge Distillation, Active Learning, Self-supervised Learning, Camera Models.
- **Programming**: PyTorch, Tensorflow, MXNet, Python, C, C++, Java, Distributed Training, Data Structure, $\LaTeX$.

Publications

- Linus Ericsson, Miguel Espinosa, **Chenhongyi Yang**, Antreas Antoniou, Amos Storkey, Shay B. Cohen, Steven McDonagh, and Elliot J. Crowley. einspace: Searching for Neural Architectures from Fundamental Operations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- **Chenhongyi Yang***, Zehui Chen*, Miguel Espinosa*, Linus Ericsson, Zhenyu Wang, Jiaming Liu, and Elliot J. Crowley. PlainMamba: Improving Non-Hierarchical Mamba in Visual Recognition. *The British Machine Vision Conference (BMVC)*, 2024.
- **Chenhongyi Yang**, Anastasia Tkach, Shreyas Hampali, Linguang Zhang, Elliot J. Crowley, and Cem Keskin. EgoPoseFormer: A Simple Baseline for Stereo Egocentric 3D Human Pose Estimation. *The European Conference on Computer Vision (ECCV)*, 2024.
- **Chenhongyi Yang**, Tianwei Lin, Lichao Huang, and Elliot J. Crowley. WidthFormer: Toward Efficient Transformer-based BEV View Transformation. *The IEEE / RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024 **Oral**.
- **Chenhongyi Yang**, Lichao Huang, and Elliot J. Crowley. Plug and Play Active Learning for Object Detection. *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*, 2024.
- Jiahao Chang, Shuo Wang, Haiming Xu, Zehui Chen, **Chenhongyi Yang**, and Feng Zhao. DETRDistill: A Universal Knowledge Distillation Framework for DETR-families. *The IEEE / CVF International Conference on Computer Vision (ICCV)*, 2023.
- **Chenhongyi Yang***, Jiarui Xu*, Shalini De Mello, and Elliot J. Crowley, Xiaolong Wang. GPViT: A High-Resolution Non-Hierarchical Vision Transformer with Group Propagation. *The International Conference on Learning Representations (ICLR)*, 2023 **Spotlight**.
- **Chenhongyi Yang**, Mateusz Ochal, Amos Storkey, and Elliot J. Crowley. Prediction-Guided Distillation for Dense Object Detection. *The European Conference on Computer Vision (ECCV)*, 2022.
- **Chenhongyi Yang**, Zehao Huang, and Naiyan Wang. QueryDet: Cascade Sparse Query for Small Object Detection. *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*, 2022 **Oral**.
- Zehui Chen*, **Chenhongyi Yang***, Qiaofei Li, Feng Zhao, Zheng-Jun Zha, and Feng Wu. Disentangle Your Dense Object Detector. *The ACM International Conference on Multimedia (ACM MM)*, 2021 **Oral**.
- Zehui Chen, **Chenhongyi Yang**, Qiaofei Li, Feng Zhao, Zheng-Jun Zha, and Feng Wu. DDOD: Dive Deeper into the Disentanglement of Object Detector. *IEEE Transactions on Multimedia*, 2023.
- Kaihong Wang, **Chenhongyi Yang**, and Margrit Betke. Consistency Regularization with High-dimensional Non-adversarial Source-guided Perturbation for Unsupervised Domain Adaptation in Segmentation. *The AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- **Chenhongyi Yang**, Vitaly Ablavsky, Kaihong Wang, Qi Feng, and Margrit Betke. Learning to Separate: Detecting Heavily-Occluded Objects in Urban Scenes. *The European Conference on Computer Vision (ECCV)*, 2020.

Preprints

- **Chenhongyi Yang**, Lichao Huang, and Elliot J. Crowley. Contrastive Object-level Pre-training with Spatial Noise Curriculum Learning. *arXiv preprint arXiv:2111.13651*, 2021.

Academic Service

- **Conference Reviewer**: CVPR, ICCV, ECCV, NeurIPS, ICLR, BMVC
- **Journal Reviewer**: TIP, TNNLS, TCSVT, TMM