

XIANG LI

xiangli12@illinois.edu | ryanxli@outlook.com | +1 2172002017 | ryanxli.github.io | Google Scholar | GitHub

EDUCATION

University of Illinois at Urbana-Champaign

Sept 2021 - Present, Urbana, IL, USA

- Ph.D. candidate in Computer Science, advised by Prof. James M. Rehg. Passed the Ph.D. qualifying exam.
- Teaching Experience: TA for Computational Photography in Spring 2023 and Fall 2023.
- Selected Coursework: Machine Learning, Computer Vision, Transfer Learning, Meta Learning, Efficient & Predictive Vision, Data Mining, Interactive Computer Graphics.

The Hong Kong University of Science and Technology

Sept 2016 - July 2020, Hong Kong

- Bachelor of Engineering in Computer Science. GPA: 3.99 / 4.3 (top 1%).
- First Class Honors, Kerry Holdings Limited Scholarship, Dean's list for all active semesters.

RESEARCH PROJECTS

Symmetry Strikes Back: From Single-Image Symmetry Detection to 3D Generation

In ArXiv, Authors: **Xiang Li**, Zixuan Huang, Anh Thai, James M. Rehg

Since Jan 2024, UIUC

- Introduces Reflect3D, a zero-shot single-image 3D reflection symmetry detector. Leveraging a scalable transformer architecture and a multi-view diffusion prior, Reflect3D achieves state-of-the-art performance in symmetry detection, robustly generalizing to in-the-wild imagery. Integrates symmetry detection into single-image 3D generation pipelines, improving structural accuracy, visual fidelity, and cohesiveness.
- Personal contribution: Formulated the problem of symmetry detection and symmetry-aware 3D generation. Designed and implemented the approach, conducted experiments, and led the paper writing process.

Video State-Changing Object Segmentation

In ICCV 2023, Authors: Jiangwei Yu*, **Xiang Li***, Xinran Zhao, Hongming Zhang, Yu-Xiong Wang

Since May 2022, UIUC

- Proposes a benchmark for video object segmentation (VOS) where the object undergoes state changes, for example, a cucumber being sliced. Reveals that current VOS models struggle with state-changing objects. Proposes an improved solution consisting of a fine-tuning strategy, representation learning, and integration of motion information.
- Personal contribution: Proposed and formulated the state change problem in the VOS setting, designed benchmark and organized data collection and annotation, and led the project in implementation, experiments, and paper writing.

YouTubePD: A Multimodal Benchmark for Parkinson's Disease Analysis

In NeurIPS Datasets and Benchmarks Track 2023, Co-first Author, Advised by Prof. Yu-Xiong Wang

Since May 2022, UIUC

- Introduces YouTubePD, the first public multimodal benchmark for Parkinson's Disease (PD) analysis, crowdsourced from existing YouTube videos featuring over 200 subjects. The benchmark provides diverse expert annotations and suggests three tasks, with experimental evaluations indicating the potential and limits of deep learning models for real-world clinical applications.
- Personal contribution: Led the project involving multiple collaborators including clinicians, led intern students in data collection, method design, and experiments, led paper writing.

FSS-1000: A 1000-Class Dataset for Few-Shot Segmentation

In CVPR 2020, Authors: **Xiang Li**, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang

Since Feb 2019, HKUST

- Highlights the key for few-shot segmentation is category variety. Proposes a few-shot segmentation dataset containing 1000 varied object categories with pixelwise annotation of ground-truth segmentation. Our proposed FSS-1000 contains a significant number of objects that have never been seen or annotated in previous datasets. Achieves comparable performance to state-of-the-art with a simple yet effective baseline method trained on our dataset.
- Personal contribution: Participated in the collection, labeling and maintenance of the dataset. Implemented the model and performed quantitative and qualitative experiments for the proposed model. Engaged in paper writing.

Cooperating RPN's: Improving Few-Shot Detection Under Distribution Shift of Object Aspect Ratios

Tech Report, Authors: **Xiang Li***, Weilin Zhang*, Yu-Xiong Wang, David A. Forsyth

Since Oct 2021, UIUC

- Proposes Cooperating RPN's (CoRPN's), an approach to improve region proposal in few-shot object detection. By training multiple distinct yet cooperating RPN's, our CoRPN's are trained to be naturally specializing in proposals of different aspect ratios, but are not overspecializing because of the cooperation constraints. Our method substantially improves few-shot detection in the large distribution shift scenario, as well as achieving state-of-the-art performance in the current benchmark.

- Personal contribution: Identified the improvements in the aspect ratio distribution shift scenario, extended a range of experiments, and revised the paper.

One-Shot Object Detection without Fine-Tuning

Tech Report, Authors: **Xiang Li***, Lin Zhang*, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang *Since June 2019, HKUST*

- As one of the first works in one-shot object detection, proposes a novel two-stage model and training strategies by integrating metric learning with an anchor-free Faster R-CNN-style detection pipeline. Eliminates the need to fine-tune on the support image and exceeds the state-of-the-art one-shot detection performance.
- Personal contribution: Proposed and implemented the ideas, performed extensive quantitative and qualitative experiments to justify our contribution, and led the paper writing.

INTERNSHIP

Student Researcher, **Google LLC**

May 2022 - Aug 2022, Seattle, US

Advised by Dr. Tianyi Zhou, Dr. Boqing Gong, and Dr. Ming-Hsuan Yang

- Investigates curriculum learning for visual continual pretraining. Designed a score function to predict model performance after continual pretraining without training to convergence. Applies the continual pretraining method to improve model performance on out-of-distribution test domains.

Intern at OpenMMLab, **SenseTime Hong Kong**

Aug 2020 - Oct 2020, Hong Kong

Advised by Dr. Kai Chen

- Developing and maintaining **MMDetection** (link), an open-source object detection framework with 11.6k stars on GitHub.
- Personal contribution: implementing detection models in MMDetection, e.g. YOLO v4; reproducing their performances and optimizing their speed, making them modular components for in-place usage of all 40 detection models; implement new functionalities for the MMDetection framework; bug fixes and code refactoring; writing and reviewing documentations; solving issues and reviewing PRs.

PROJECTS

- **CUDA CNN Forward** (*Sept 2018 - Dec 2018*) Re-implemented the CNN forward algorithm of MXNet in CUDA using im2col + GEMM among other optimizations. UIUC Parallel Programming final project, performance ranked top 10 out of over 60 groups.
- **System Programming** (*Sept 2018 - Dec 2018*) Implemented malloc, parallel make, shell, a command-line text editor, and a server-client model in C. UIUC System Programming course projects.
- **Computer Graphics** (*Jan 2018 - May 2018*) Implemented a Ray Tracer and an interactive image processor in C++, a Modeler and an Animator with OpenGL. HKUST Computer Graphics course projects.

ACADEMIC SERVICES AS REVIEWER

- 2024: CVPR, ECCV, NeurIPS, NeurIPS dataset and benchmark track.
- 2023: CVPR, ICCV, NeurIPS, NeurIPS dataset and benchmark track.
- 2022: CVPR, ECCV, NeurIPS, NeurIPS dataset and benchmark track.
- 2021: CVPR, ICCV.

SKILLS

- Programming Languages: Python, C/C++, CUDA, Java, HTML + CSS + JavaScript.
- Libraries: PyTorch, TensorFlow, JAX/Flax, Scikit-learn, Numpy, OpenCV, PIL, OpenGL, WebGL, React.
- Languages: English (Fluent), Mandarin Chinese (Native).