

AN YAN

✉ ayan@ucsd.edu · ☎ (+1) 858-275-4651 · in An Yan

🎓 EDUCATION

University of California, San Diego	La Jolla, CA
Ph.D. in Computer Science	2020 – 2024
M.S. in Computer Science	2018 – 2020
Advisor: Julian McAuley	
University of Science and Technology of China	Hefei, China
B.E. in Electronic Engineering and Information Science	2014 – 2018

⚙️ TECH SKILLS

- **Programming Languages:** Python, C, SQL, MATLAB, Haskell
- **Tools & Frameworks:** PyTorch, HuggingFace, PySpark, Hive, Azure, AWS, Distributed training, Linux

👤 INDUSTRY EXPERIENCE

Research Intern at Microsoft , Redmond, WA	Sep. 2023 – Present
Hosts: Zhengyuan Yang, Jianfeng Wang, Linjie Li, Kevin Lin, Zicheng Liu, Lijuan Wang	
<ul style="list-style-type: none">• Training, evaluation and understanding of Large Vision-Language Models.• Working on new applications with GPT-4V.	
Research Intern at Adobe , San Jose, CA	June. 2023 – Sep. 2023
Hosts: Raghav Addanki, David Arbour, Zhao Song, Tong Yu	
<ul style="list-style-type: none">• Gradient-based Constrained sampling from LMs (GPT-2, Flan-T5, Llama) to optimize user engagement.• Proposed a method to generate new promotion text with 15% user score improvement over state-of-the-arts.	
Research Intern at Meta , Menlo Park, CA	June. 2022 – Sep. 2022
Hosts: Cem Akkaya, Licheng Yu, Charlie Zhu, Yang Bai	
<ul style="list-style-type: none">• Multi-modal pre-training for ads understanding and generation, on 50+ million ads from Facebook & IG.• Prototyped an extract-and-generate framework to customize BART to generate template ads for advertisers.	
Applied Scientist Intern at Amazon , Seattle, WA	June. 2021 – Sep. 2021
Hosts: Chaosheng Dong, Yan Gao, Jinmiao Fu, Tong Zhao	
<ul style="list-style-type: none">• Personalized complementary recommendation. Built a transformer and GNN framework with contrastive augmentation for user modeling and recommendation. 25% relative improvement on hit rate over baselines.• Paper at WWW-2022. Top 10 most viewed publications of 2022 at Amazon Science: https://www.amazon.science/latest-news/the-most-viewed-amazon-science-publications-of-2022	
Applied Scientist Intern at Amazon , Santa Barbara, CA	June. 2020 – Sep. 2020
Hosts: Craig Bennett, Nic Jedema	
<ul style="list-style-type: none">• Question-Answering quality evaluation. Collected and cleaned 1 million internal QA pairs via Spark.• Finetuned BERT models for evaluating QA quality at Alexa.	

⚙️ SELECTED PROJECTS

Complex Contextual Recommendation

- Collaboration with Google Deepmind. Proposed a new task for LLMs & Recommendation. Generating datasets by prompting ChatGPT. Training LM adapters for item retrieval.

Robust and Interpretable Visual Understanding via Natural Language

- Proposed a new paradigm for visual recognition via pruned attributes generated from LLMs.
- Papers at ICCV-2023, MedNeurIPS-2023, WACV-2024.

Personalized Multi-modal Explanation Generation

- Collected a large scale web image-text dataset on google local. Proposed a new task and visual-GPT framework for generating personalized multi-modal explanations for recommendation. Paper at SIGIR-2023.

Language Models for Clinical NLP

- Finetuning Language Models for Clinical NLP. Built popular Medical LMs (RadBERT) on HuggingFace with 1K downloads per month and 100K+ peak monthly downloads. Paper at Journal of Radiology.

PUBLICATIONS

Preprints & Under submission:

GPT-4V in Wonderland: Large Multimodal Models for Zero-Shot Smartphone GUI Navigation

- **An Yan**, Zhengyuan Yang, Wanrong Zhu, Kevin Lin, Linjie Li, Jianfeng Wang, Jianwei Yang, Yiwu Zhong, Julian McAuley, Jianfeng Gao, Zicheng Liu, Lijuan Wang

GPT-4V as a Generalist Evaluator for Vision-Language Tasks

- Xinlu Zhang, Yujie Lu, Weizhi Wang, **An Yan**, Jun Yan, Lianke Qin, Heng Wang, Xifeng Yan, William Yang Wang, Linda Ruth Petzold

CLIP also Understands Text: Prompting CLIP for Phrase Understanding

- **An Yan**, Jiacheng Li, Wanrong Zhu, Yujie Lu, William Yang Wang, Julian McAuley

Refereed Publications:

Driving through the Concept Gridlock: Unraveling Explainability Bottlenecks in Automated Driving

- Jessica Echterhoff, **An Yan**, Kyungtae Han, Amr Abdelraouf, Rohit Gupta, Julian McAuley
- Winter Conference on Applications of Computer Vision (**WACV-2024**)

Mitigating Spurious Correlations for Medical Image Classification via Natural Language Concepts

- **An Yan**, Yu Wang, Petros Karypis, Zexue He, Amilcare Gentili, Chun-Nan Hsu, Julian McAuley
- Conference on Neural Information Processing Systems (**NeurIPS 2023**) Medical Imaging Workshop

M4: A Multi-Level, Multi-Task, and Multi-Domain Medical Benchmark for Language Model Evaluation

- Zexue He, Yu Wang, **An Yan**, Yao Liu, Eric Y Chang, Amilcare Gentili, Julian McAuley, Chun-Nan Hsu
- Empirical Methods in Natural Language Processing (**EMNLP-2023**)

Learning Concise and Descriptive Attributes for Visual Recognition

- **An Yan**, Yu Wang, Yiwu Zhong, Chengyu Dong, Zexue He, Yujie Lu, William Wang, Jingbo Shang, Julian McAuley
- International Conference on Computer Vision 2023 (**ICCV-2023**)

Personalized Showcases: Generating Multi-Modal Explanations for Recommendations

- **An Yan**, Zhankui He, Jiacheng Li, Tianyang Zhang, Julian McAuley
- The International ACM SIGIR Conference on Research and Development in Information Retrieval (**SIGIR-2023**)

Visualize Before You Write: Imagination-Guided Open-Ended Text Generation

- Wanrong Zhu, **An Yan**, Yujie Lu, Wenda Xu, Xin Eric Wang, Miguel Eckstein, William Yang Wang
- European Chapter of the Association for Computational Linguistics (**EACL-2023**)

ImagineE: An Imagination-Based Automatic Evaluation Metric for Natural Language Generation

- Wanrong Zhu, Xin Wang, **An Yan**, Miguel Eckstein, William Wang
- European Chapter of the Association for Computational Linguistics (**EACL-2023**)

Disambiguating Medical Reports via Contrastive Knowledge Infusion

- Zexue He, **An Yan**, Amilcare Gentili, Julian McAuley, Chun-Nan Hsu
- AAAI Conference on Artificial Intelligence (**AAAI-2023**)

Robust multi-view fracture detection in the presence of other abnormalities using HAMIL-Net

- Xing Lu, Eric Chang, Jiang Du, **An Yan**, Julian McAuley, Amilcare Gentili, Chunnan Hsu
- Military Medicine, 2023

RadBERT: Adapting Language Models to Radiology

- **An Yan**, Chun-Nan Hsu, Amilcare Gentili, Julian McAuley
- Journal of Radiology: Artificial Intelligence, 2022

Personalized Complementary Product Recommendation

- **An Yan**, Yan Gao, Chaosheng Dong, Jinmiao Fu, Tong Zhao, Yi Sun, Julian McAuley
- The ACM Web Conference (WWW-2022)

Semi-supervised Multi-Label Classification with 3D CBAM Resnet for Tuberculosis Cavern Report

- Xing Lu, **An Yan**, Eric Y Chang, Chun-nan Hsu, Julian McAuley, Jiang Du, Amilcare Gentili
- Conference and Labs of the Evaluation Forum (CLEF-2022)

Weakly Supervised Contrastive Learning for Chest X-Ray Report Generation

- **An Yan**, Zexue He, Xing Lu, Jiang Du, Eric Chang, Amilcare Gentili, Julian McAuley, Chun-Nan Hsu
- Empirical Methods in Natural Language Processing (EMNLP-2021), Findings

Describing Visual Differences Needs Semantic Understanding of Individuals

- **An Yan**, Xin Wang, Tsu-Jui Fu, William Wang
- European Chapter of the Association for Computational Linguistics (EACL-2021)

Multimodal Style Transfer Learning for Outdoor Vision-and-Language Navigation

- Wanrong Zhu, Xin Wang, Tsu-Jui Fu, **An Yan**, Pradyumna Narayana, Kazoo Sone, Sugato Basu, William Wang
- European Chapter of the Association for Computational Linguistics (EACL-2021)

2D Convolutional Neural Networks for Sequential Recommendation

- **An Yan**, Shuo Cheng, Wang-Cheng Kang, Mengting Wan, Julian McAuley
- ACM International Conference on Information and Knowledge Management (CIKM-2019)

PA3D: Pose-Action 3D Machine for Video Recognition

- **An Yan**, Yali Wang, Zhifeng Li, Yu Qiao
- IEEE Conference on Computer Vision and Pattern Recognition (CVPR-2019)

ACADEMIC EXPERIENCE

Research Assistant at UC San Diego, La Jolla, CA 2018 – Present

Advisor: Prof. Julian McAuley

- Vision and Language. Personalization and Recommendation. (w/ Toyota, Google Deepmind)

Visiting Student at UC Santa Barbara, Goleta, CA June. 2019 – June. 2020

Advisor: Prof. William Wang

- Vision-Language Navigation (w/ ByteDance AI Lab, Google research)
- Image Captioning, Visual Machine Translation.

Visiting Student at Chinese Academy of Sciences, Shenzhen, China Mar. 2018 – Aug. 2018

Advisor: Prof. Yu Qiao

- Human Action understanding, Human Pose Estimation. (w/ Tencent AI Lab)

Research Assistant at USTC, Hefei, China Sep. 2016 – Jan. 2018

Advisors: Prof. Cong Shen, Prof. Weiping Li

- Reinforcement Learning, Multi-Armed Bandits.

PROFESSIONAL SERVICES

- Reviewer: NeurIPS, ICLR, ACL, CIKM, EMNLP, AAAI, RecSys, WWW, WSDM, MLHC, EACL
- Senior PC member: AAAI-2023, AAAI-2024
- Session Chair: WWW-2022