Yun-Cheng (Joe) Wang

Mountain View, CA | joewang622@gmail.com | Google Scholar | joewang622

yunchengwang.github.io

About Me

I am a **machine learning** researcher and engineer devoted to developing **efficient**, **lightweight**, and **scalable** systems, with expertise in **knowledge graphs (KGs)**, **natural language processing (NLP)**, **large language models (LLMs)**, **information retrieval**, and **on-device AI**.

With a strong commitment to impactful research, I have led collaborations across industry and academia, developing AIML solutions that serve **millions of users**. I am also dedicated to mentoring and fostering innovation in teams, driving forward both individual growth and **cutting-edge AIML technologies**.

Education

University of Southern California – Los Angeles, CA Jan 2021 - Dec 2023
Ph.D. in Electrical Engineering

- Dissertation: Green Knowledge Graph Completion and Scalable Generative Content Delivery

University of Southern California – Los Angeles, CA Aug 2018 - Dec 2019
M.S. in Electrical Engineering

- Relevant Coursework: Pattern Recognition, Multimedia Compression, Convex Optimization

National Taiwan University – Taipei, Taiwan Sep 2014 - Jun 2018
B.S. in Electrical Engineering

- Relevant Coursework: Digital Speech Processing, Machine Learning Foundations, Artificial Intelligence

Professional Experiences

Research Scientist, Yahoo, Inc. – Mountain View, CA Jan 2024 – Present

- Developed machine learning solutions to support the query intent service (QIS) in Yahoo Search.
- The models serve more than 900M monthly active users with around 27% reduction in latency.
- Collaborated with product teams to apply machine learning models to various use cases.

Research Intern, Yahoo, Inc. – Remote Jun 2023 – Aug 2023

- Innovated a fact ranking mechanism to generate knowledge-grounded entity descriptions.
- Curated a high-quality data-to-text dataset containing 20K examples using LLMs and KG fact-checking.

Data Science Intern, Taboola, Inc. – Los Angeles, CA Jun 2019 – Aug 2019

- Discovered trending topics in news articles through network analysis.
- The topic graph was incrementally updated based on over 20K daily articles.

Publications

Conference Papers

[1] Tsung-Shan Yang, **Yun-Cheng Wang**, Chengwei Wei, C.-C. Jay Kuo, "GHOI: A Green Human-Object-Interaction Detector," *IEEE International Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2024.

[2] **Yun-Cheng Wang**, Xiou Ge, Bin Wang, C.-C. Jay Kuo, "AsyncET: Asynchronous Representation Learning for Knowledge Graph Entity Typing," *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2024.

[3] **Yun-Cheng Wang**, Xiou Ge, Bin Wang, C.-C. Jay Kuo, "GreenKGC: A Lightweight Knowledge Graph Completion Method," *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.

[4] Xiou Ge, **Yun-Cheng Wang**, Bin Wang, C.-C. Jay Kuo, "Compounding Geometric Operations for Knowledge Graph Completion," *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.

[5] Zhanxuan Mei, **Yun-Cheng Wang**, Xingze He, C.-C. Jay Kuo, "GreenBIQA: A Lightweight Blind Image Quality Assessment Method," *IEEE International Workshop on Multimedia Signal Processing (MMSP)*, 2022. **(Top 10% Paper Award)**

[6] Bin Wang, Fenxiao Chen, **Yun-Cheng Wang**, C.-C. Jay Kuo, "Efficient Sentence Embedding via Semantic Subspace Analysis," *IEEE International Conference on Pattern Recognition (ICPR)*, 2020.

Journal Papers

[1] Xiou Ge, **Yun-Cheng Wang**, Bin Wang, C.-C. Jay Kuo, "Knowledge Graph Embedding: An Overview," *APSIPA Transactions on Signal and Information Processing*, 2024.

[2] Jintang Xue, **Yun-Cheng Wang**, Chengwei Wei, Xiaofeng Liu, Jonghye Woo, C.-C. Jay Kuo, "Bias and Fairness in Chatbots: An Overview," *APSIPA Transactions on Signal and Information Processing*, 2024.

[3] Xiou Ge, **Yun-Cheng Wang**, Bin Wang, C.-C. Jay Kuo, "Knowledge Graph Embedding with 3D Compound Geometric Transformations," *APSIPA Transactions on Signal and Information Processing*, 2024.

[4] Zhanxuan Mei*, **Yun-Cheng Wang***, Xingze He, Yong Yan, C.-C. Jay Kuo, "Lightweight High-Performance Blind Image Quality Assessment," *APSIPA Transactions on Signal and Information Processing*, 2024.

[5] **Yun-Cheng Wang**, Jintang Xue, Chengwei Wei, C.-C. Jay Kuo, "An Overview on Generative AI at Scale with Edge-Cloud Computing," *IEEE Open Journal of the Communications Society*, 2023.

[6] Chengwei Wei, **Yun-Cheng Wang**, Bin Wang, C.-C. Jay Kuo, "An Overview on Language Models: Recent Developments and Outlook," *APSIPA Transactions on Signal and Information Processing*, 2023.

[7] Xiou Ge, **Yun-Cheng Wang**, Bin Wang, C.-C. Jay Kuo, "TypeEA: Type-Associated Embedding for Knowledge Graph Entity Alignment," *APSIPA Transactions on Signal and Information Processing*, 2023.

[8] **Yun-Cheng Wang**, Xiou Ge, Bin Wang, C.-C. Jay Kuo, "KGBoost: A Classification-Based Knowledge Base Completion Method with Negative Sampling," *Pattern Recognition Letters*, 2022.

[9] Xiou Ge, **Yun-Cheng Wang**, Bin Wang, C.-C. Jay Kuo, "CORE: A Knowledge Graph Entity Type Prediction Method via Complex Space Regression and Embedding," *Pattern Recognition Letters*, 2022.

[10] Fenxiao Chen, **Yun-Cheng Wang**, Bin Wang, C.-C. Jay Kuo, "Graph Representation Learning: A Survey," *APSIPA Transactions on Signal and Information Processing*, 2020.

[11] Bin Wang, Angela Wang, Fenxiao Chen, **Yun-Cheng Wang**, C.-C. Jay Kuo, "Evaluating Word Embedding Models: Methods and Experimental Results," *APSIPA Transactions on Signal and Information Processing*, 2019. **(2022 Sadaoki Furui Prize Paper Award)**

Preprints

[1] Tsung-Shan Yang, **Yun-Cheng Wang**, Chengwei Wei, Suya You, C.-C. Jay Kuo, "Efficient Human-Object-Interaction (EHOI) Detection via Interaction Label Coding and Conditional Decision," *arXiv preprint*, 2024.

[2] Jintang Xue, **Yun-Cheng Wang**, Chengwei Wei, Xiaofeng Liu, Jonghye Woo, C.-C. Jay Kuo, "Word Embedding Dimension Reduction via Weakly-Supervised Feature Selection," *arXiv preprint*, 2024.

[3] Zhanxuan Mei, **Yun-Cheng Wang**, C.-C. Jay Kuo, "GSBIQA: Green Saliency-guided Blind Image Quality Assessment Method," *arXiv preprint*, 2024.

[4] Zhanxuan Mei, **Yun-Cheng Wang**, C.-C. Jay Kuo, "GreenSaliency: A Lightweight and Efficient Image Saliency Detection Method," *arXiv preprint*, 2024.

[5] Zhanxuan Mei, **Yun-Cheng Wang**, C.-C. Jay Kuo, "Blind Video Quality Assessment at the Edge," *arXiv preprint*, 2023.

Invited Talks

- "Scalable Generative Content Delivery", Research Center for Information Technology Innovation (CITI) at Academia Sinica, Jan 2024.
- "Green Learning with Knowledge Graphs", National Tsing Hua University (NTHU), Dec 2023.
- "Knowledge Graph: Foundations and Applications", National Taiwan Normal University (NTNU), Sep 2020.

Research Projects

Decoupling Semantics and Syntax in Language Models Aug 2023 - Present

- Developed lightweight and trustworthy models for domain-specific language generation, e.g., bio-medical.
- Adopted knowledge graphs to capture semantic patterns in natural language.
- Modualized language models through a bottom-up manner to achieve efficiency and interpretability.

Multi-modality Alignment, *Sponsored by Army Research Lab (ARL)* Aug 2023 - Present

- Leveraged the embedding space to connect different modalities for multi-modal reasoning.
- Extracted human-object interactions using spatial and latent features with hierarchical classifiers.
- Devised an alignment module in a joint embedding space for text-to-image and image-to-text retrieval.

Scalable Generative AI Services under Edge-Cloud Computing Jan 2023 - Oct 2023

- Analyzed the memory, computation, and network requirements to deploy GenAI services, e.g., ChatGPT, across different scales.
- Estimated the latency for GenAI services under different communication frameworks.
- Identified considerations when designing GenAI systems with better efficiency, computation offloading, and privacy.

Efficient Reasoning on KGs using Lightweight Models Jan 2021 - Oct 2023

- This project aimed at predicting missing information, including entity types and relations, in knowledge graphs using lightweight models.
- Leveraged feature pruning to achieve parameter efficiency and SOTA performance in low dimensions.
- Proposed novel modeling of entity types to improve expressiveness while retaining scalability to large KGs.
- Innovated an asynchronous KGE learning framework to improve performance on both link prediction and entity type prediction tasks.
- Overall, inference FLOPs were reduced 100 times, and the number of parameters was reduced 15 times.

Build Image and Video Quality Assessment, *Sponsored by Meta* Aug 2021 - Dec 2022

- Predicted user perceptual experience on images and videos without references.
- Developed a lightweight pipeline that can be deployed on user devices and manage real-time inference.
- Achieved SOTA results on synthetic datasets and competitive results on authentic datasets with a 54 times smaller model size.

Teaching Experiences

Guest Lecturer - University of Southern California Fall 2022, Fall 2023

- Course: Multimedia Compression
- Delivered a lecture on "Blind Image Quality Assessment"

Teaching Assistant - University of Southern California Spring 2023 - Fall 2023

- Course: Stochastic Process for Financial Engineering
- Facilitated course delivery and provided guidance through discussions and office hours

Course Mentor - University of Southern California Fall 2019

- Course: Probability for Electrical and Computer Engineers
- Supported course design and hosted office hours

Academic Services

Conference Reviewer

- International Conference on Knowledge Discovery and Data Mining (KDD): 2024, 2025
- ACL Rolling Review (ARR): Feb 2024, Apr 2024, Jun 2024, Aug 2024
- Empirical Methods in Natural Language Processing (EMNLP): 2023
- Asian Conference on Machine Learning (ACML): 2024
- European Conference on Machine Learning (ECML): 2022

Journal Reviewer

- IEEE/ACM Transactions on Audio, Speech and Language Processing (T-ASL), 2 times
- IEEE Internet of Things Magazine (IoTM), 2 times
- IEEE Journal of Biomedical and Health Informatics (JBHI), 2 times
- IEEE Transactions on Artificial Intelligence (TAI)
- APSIPA Transactions on Signal and Information Processing, 7 times
- Applied Sciences (MDPI)
- Buildings (MDPI)
- Information (MDPI), 2 times
- Knowledge and Information Systems (Springer Nature), 2 times
- Neural Processing Letters (Springer Nature)

Professional Affiliations

- Member, ACM Special Interest Group on Knowledge Discovery and Data Mining (SIGKDD)
- Member, Association for Computational Linguistics (ACL)
- Member, IEEE Signal Processing Society

Awards

- USC Viterbi Graduate Fellowship/Research Assistantship/Teaching Assistantship — Jan 2021 - Dec 2023
- Sadaoki Furui Prize Paper Award (Pioneering Contributions in Speech Processing) — APSIPA ASC 2022
- Top 10% Paper Award — IEEE MMSP 2022
- USC EE MS honors program — Aug 2019 - Dec 2019

Technical Skills

Programming Languages: Python, C++, Java, SQL, SPARQL, Bash

Big Data & ML: Hadoop, PySpark, PyTorch, TensorFlow, XGBoost, LightGBM, FastText, Transformers, Pandas