MAN LUO

+1 (480) 869-7882 ⬦ luoman.cs@gmail.com ⬦ Website ⬦ Google Scholar

PROFESSIONAL SUMMARY

Applied AI researcher with 6+ years of experience in agentic AI, multimodal retrieval and generation (RAG), synthetic data generation pipelines, foundation vision-language model training and evaluation. Industry experience includes deploying web agents, retrieval-augmented generation (RAG) system, and AI for healthcare at Intel, Mayo Clinic, Google, Meta, and Salesforce. Published in top venues including ICML, ACL, EMNLP, AAAI, NAACL, and WACV.

EDUCATION

Ph.D., Computer Science   *August 2018 - May 2023*
Arizona State University (ASU), Tempe, Arizona, USA
Thesis: Neural Retriever-Reader for Information Retrieval and Question Answering
Chair: Dr. Chitta Baral
Committee: Dr. Yezhou Yang, Dr. Eduardo Blanco, Dr. Danqi Chen

Bachelor of Science, Computer Science   *September 2014 - July 2018*
Beijing Forestry University, Beijing, China

RESEARCH AREA

Multimodal and language models post-training, Multimodal Retrieval and Generation, Synthetic Data Pipelines, Computer Use Agents, and Multimodal Applications in Healthcare.

INDUSTRY RESEARCH EXPERIENCE

Senior AI Research Scientist at Multimodal Cognitive Team, Intel Lab   *March 2024 - Present*
Developed agentic computer use agents using reinforcement learning (DPO, GRPO with TRL), scalable synthetic data pipelines via prompt engineering, and post-training methods for foundation VLMs (2-70B) on tasks like VQA, retrieval, and instruction following. Designed diffusion-based image generation pipeline with embedded text rendering.

Research Fellow at Mayo Clinic   *June 2023 - March 2024*
Built VLM and LLM pipelines for radiology report generation and longitudinal clinical note extraction using private EHR data, leveraging multimodal modeling and prompt-based fine-tuning.

Research Intern at Google Research   *Aug 2022 - Dec 2022*
Designed retrieval-augmented in-context learning systems using DPR-like and BM25 retrievers to improve LLM performance on few-shot QA and multi-hop reasoning benchmarks through effective context selection and knowledge grounding.

Research Intern at Meta Reality Lab   *May 2022 - Aug 2022*
Built memory-efficient hybrid retrievers and adversarial robustness evaluation frameworks for AR/VR retrieval-augmented generation (RAG) applications, combining dense-sparse retrieval and contrastive perturbation testing.

Research Intern at Salesforce.Inc   *May 2021 - Aug 2021*
Evaluated generative vs. extractive QA models under domain shifts, proposing a unified benchmarking framework and deployment recommendations to guide model selection and robustness in real-world settings.

TEACHING/MENTORING

ASU Capstone Project Mentor
Detect and rewrite the toxicity in paper reviews, 5 students,   *Sep 2023 - Present*
Ph.D. Mentor

Xiang Rui (Ph.D. Student at Arizona State University).	*Aug 2024 - Present*
Md Messal Monem Miah (Ph.D. Student at Texas A&M University).	*Oct 2023 - Present*

Master Mentor

Sanyam Lakhanpal (Master Student at ASU).	*Oct 2023 - Apirl 2024*
Shrinidhi Kumbhar (Master Student at ASU).	*Jan 2023 - June 2023*
Srija Macherla (SWE at Amazon).	*Jan 2022 - Jun 2022*
Yankai Zeng (Ph.D student at The University of Texas at Dallas).	*Aug 2020 - June 2021*

NLP Course Project Mentor

Domain Oriented Question Generation, 26 students,	*Aug 2021 - Dec 2021*
Differential Diagnosis Dialogue Generation, 20 students,	*Aug 2021 - Dec 2021*
Semantic Information Availability (SIA) Task, 5 students,	*Jun 2020 - May 2020*
Question Answering with Varied Types of Reasoning, 5 students.	*Jun 2020 - May 2020*

Teaching Assistant

CSE259 Logic in Computer Science	*Dec 2020 - Dec 2021*
CSE579 Knowledge Representation and Reasoning	*Aug 2019 - Dec 2019*
CSE205 Object-Oriented Programming and Data Structures	*Aug 2018 - Dec 2018*

ACADEMIC SERVICE

Finance and Sponsorship Chair, WiML Workshop at NeurIPS 2025.
Editor, PLOS Digital Health.
Long-Term Super-Volunteers, WiML Workshop at NeurIPS 2024.
Organizer, Multimodal4Health at ICHI 2024
Guest Editor, PLOS Digital Medicine.
Organizer, O-DRUM at CVPR 2023.
Organizer, O-DRUM at CVPR 2022.
Reviewers, Nature Machine Intelligent, ACL, NAACL, EMNLP, EACL, AAAI, Neurips, IROS.

INVITED TALK

"Inspecting the Rise of Multimodal Through Retrieval and Content Generation Tasks" at UIUC	*Oct 2024*
"Synthetic Data for Generalization and Efficiency" at ASU	*Sep 2024*
"Retrieval Based In-context Learning for Large Language Models" at Google	*Mar 2024*
"Advancing Multimodal Retrieval and Generation" at UMBC	*Dev 2023*
"Transformer-based Multimodal Generative Model" at Mayo Clinic Radiology Showcase	*Nov 2023*
"The Trend of Transformer-based Multimodal in Radiology" at RSNA	*Nov 2023*
"Visual-Retriever-Reader for Knowledge-based Question Answering" at SERUM WACV	*Jan 2023*
"Semantic Searching in Biomedical Domain" at exploreCSR workshop (ASU).	*Mar 2021*

AWARD

Finalist of 2021 Knowledge Mobilization Awards. Website	*April 2021*
2019 ICLP conference Doctoral Consortium Travel Award. Website	*September 2019*
Honorable Mention in Interdisciplinary Contest in Modeling(ICM)	*April 2017*

PUBLICATION

- Su, X., **Luo, M.**, Pan, K., Chou, T., Lal, V., Howard, P. SK-VQA: Synthetic Knowledge Generation at Scale for Training Context-Augmented Multimodal LLMs. ICML 2025 (Selected as Oral, 1%) .

- Zhang, P., Jin, H., Hu, L., Li, X., Kang, L., **Luo, M.**, Song, Y. and Wang, H REVOLVE: Optimizing AI Systems by Tracking Response Evolution in Textual Optimization. ICML 2025.

- U Ghaffar, A Tariq, M Choudry, L Briggs, A Channar, I Banerjee, **M Luo** Domain-specific large language model for predicting prostate cancer treatment plan Journal of Clinical Oncology 2025.

- **Luo, M.** Warren, C., Cheng, Lu., Abdul-Muhsin, H., Banerjee, I. Assessing Empathy in Large Language Models with Real-World Physician-Patient Interactions. IEEE BigData 2024.

- Lakhanpal, S., Chopra, S., Jain, V., Chadha, A., **Luo, M.** Refining Text-to-Image Generation: Towards Accurate Training-Free Glyph-Enhanced Image Generation. WACV 2025.

- **Luo, M.**, Xu, X., Liu, Y., Pasupat, P., Kazemi. In-context Learning with Retrieved Demonstrations for Language Models: A Survey. TMLR Journal 2024.

- **Luo, M. et al** Automated Extraction of Patient-Centered Outcomes following Breast Cancer Treatment: An Open-Source Large Language Model-Based Toolkit. JCO Clinical Cancer Informatics 2024.

- Parmar, M., Patel, N., Varshney, N., Nakamura, M., **Luo, M.**, Mashetty, S., Mitra, A., Baral, C. Towards Systematic Evaluation of Logical Reasoning Ability of Large Language Models. ACL 2024.

- Chiang, C. C., **Luo, M.**, Dumkrieger, G., Trivedi, S., Chen, Y. C., Chao, C. J., ... & Banerjee, I. A Large Language Model-Based Generative Natural Language Processing Framework Finetuned on Clinical Notes Accurately Extracts Headache Frequency from Electronic Health Records. Headache: The Journal of Head and Face Pain 2024.

- **Luo, M.**, Xu, X., Dai, Z., Pasupat, P., Kazemi, M., Baral, C., ... Zhao, V. Y. Dr. ICL: Demonstration-Retrieved In-context Learning. NeurIPS 2023 Workshop R0-FoMo.

- **Luo, M.**, Tariq, A., Patel, B., Banerjee, I. M3-X: Multimodal Generative Model for Screening Mammogram Reading and Explanation Medical Imaging Meets NeurIPS 2023.

- Varshney, N., **Luo, M.**, Baral, C. Exploring Training Objectives for Passage-level Differentiable Search Indexing SocalNLP 2023.

- **Luo, M.**, Tariq, A., Patel, B., Banerjee, I. Transformer-based Multimodal Generative Model: Use-case of Screening Mammogram Reading. RSNA 2023.

- **Luo, M.** Fang, Z. Gokhale, T. Baral, C. End-to-end Knowledge Retrieval with Multi-modal Queries. ACL 2023.

- **Luo, M.**, Jain, S., Gupta, A., Einolghozati, A., Oguz, B., Chatterjee, D., Chen, X., Baral, C. and Heidari, P., 2022. A Study on the Efficiency and Generalization of Light Hybrid Retrievers. ACL 2023.

- Parmar, M., Mishra, S., Purohit, M., **Luo, M.**, Baral, C. In-BoXBART: Get Instructions into Biomedical Multi-Task Learning. NAACL 2022 Findings.

- Gokhale, T., Mishra, S., **Luo, M.**, Sachdeva, B., Baral, C. Generalized but not Robust? Comparing the Effects of Data Modification Methods on Out-of-Domain Generalization and Adversarial Robustness. ACL 2022 Findings.

- **Luo, M.**, Mitra, A., Gokhale, T., Baral, C. Improving Biomedical Information Retrieval with Neural Retrievers. AAAI 2022.

- **Luo, M.**, Zeng, Y., Banerjee, P., Baral, C. Weakly-Supervised Visual-Retriever-Reader for Knowledge-based Question Answering. EMNLP 2021.

- **Luo, M.** Sampat, S. Tallamn, R. Zeng, Y. Vancha, M. Sajja, A. Baral, C. Just because you are right, doesnt mean I am wrong: Overcoming a bottleneck in development and evaluation of Open-Ended VQA tasks. EACL 2021.

- Lee, J. and **Luo, M.**, 2019. Strong equivalence for LPMLN programs. ICLP 2019.

- Varshney, N., **Luo, M.**, Baral, C. Can Open-Domain QA Reader Utilize External Knowledge Efficiently like Humans? AAAI 2023 Workshop on Knowledge Augmented Methods for NLP

- **Luo, M.**, Parmar, M., Mahendran, J. S., Jain, S., Rawal, S., Baral, C. SCONER: Scoring Negative Candidates Before Training Neural Re-Ranker For Question Answering ICML 2022 Workshop on Knowledge Retrieval and Language Models.

- **Luo, M.**, Saxena, S., Mishra, S., Parmar, M., Baral, C. BioTABQA: Instruction Learning for Biomedical Table Question Answering CEUR Workshop 2022.

- **Luo, M.** Neural Retriever and Go Beyond: A Thesis Proposal. NAACL 2022 Student Research Workshop.

- **Luo, M.**, Chen, S., Baral, C A Simple Approach to Jointly Rank Passages and Select Relevant Sentences in the OBQA Context NAACL 2022 Student Research Workshop.

- **Luo, M.**, Hashimoto, K., Yavuz, S., Liu, Z., Baral, C., Zhou, Y. Choose Your QA Model Wisely: A Systematic Study of Generative and Extractive Readers for Question Answering ACL 2022 Spa-NLP workshop.

PRE-PRINT

- ML Olson, N Ratzlaff, M Hinck, **M Luo**, S Yu, C Xue, V Lal Semantic Specialization in MoE Appears with Scale: A Study of DeepSeek R1 Expert Specialization arXiv preprint 2025.

- **M Luo**, B Peterson, R Gan, H Ramalingame, N Gangrade, A Dimarogona I Banerjee, P Howard Benchmark on Peer Review Toxic Detection: A Challenging Task with a New Dataset arXiv preprint 2025.

- E Aflalo, GBM Stan, T Le, **M Luo**, S Rosenman, S Paul, SY Tseng, V Lal FiVL: A Framework for Improved Vision-Language Alignment arXiv preprint 2025.

- N Ratzlaff, **M Luo**, X Su, V Lal, P Howard Training-Free Mitigation of Language Reasoning Degradation After Multimodal Instruction Tuning arXiv preprint 2025.

- G Ben-Melech Stan, E Aflalo, **M Luo**, S Rosenman, T Le, S Paul, SY Tseng, V Lal FastRM: An efficient and automatic explainability framework for multimodal generative models arXiv preprint 2025.

- PK Choubey, X Su, **M Luo**, X Peng, C Xiong, T Le, S Rosenman, V Lal, P Mui, R Ho, P Howard, CS Wu Distill-SynthKG: Distilling Knowledge Graph Synthesis Workflow for Improved Coverage and Efficiency arXiv preprint 2024.

- **Luo, M.**, Kumbhar, S., Parmar, M., Varshney, N., Banerjee, P., Aditya, S., Baral, C. Towards LogiGLUE: A Brief Survey and A Benchmark for Analyzing Logical Reasoning Capabilities of Language Models. arXiv preprint 2023.

- Macherla, S., **Luo, M.**, Parmar, M., Baral, C. MDDial: A Multi-turn Differential Diagnosis Dialogue Dataset with Reliability Evaluation. arXiv preprint 2023.

- **Luo, M.**. Neural Retriever-Reader for Information Retrieval and Question Answering (Doctoral dissertation, Arizona State University, 2023).

- Varshney, N., Parmar, M., Patel, N., Handa, D., Sarkar, S., **Luo, M.**, Baral, C.. Can NLP Models Correctly Reason Over Contexts that Break the Common Assumptions?. arXiv preprint 2023.

- Liu, Z., Chen, Y., Li, J., **Luo, M.**, Yu, P. S., Xiong, C. Improving contrastive learning with model augmentation. arXiv preprint 2022.

- Banerjee, P., Baral, C., **Luo, M.**, Mitra, A., Pal, K., Son, T. C., Varshney, N. Can Transformers Reason About Effects of Actions? arXiv preprint, 2020.

BOOK MANUSCRIPT

Luo, M. Gokhale, T., Varshney, N., Yang, Y., Baral, C. Advances in Multi-Modal Information Retrieval. *Springer Nature*