Qiao Jin, M.D.
qiao.jin@nih.gov
www.qiaojin.info
Research Fellow
National Library of Medicine
National Institutes of Health

I am an AI researcher working on language modeling for biomedicine. My long-term goal is to democratize biomedical knowledge by building AI systems that provide accurate, verifiable, and understandable information to physicians, patients, and scientists. Currently, I work on medical AI evaluation, retrieval-augmented generation, language agents, and AI for clinical trials.

Education

2014–2022 **B.Sc.** (2019), **M.D.** (2022), Tsinghua University
Thesis: *Large-scale Text Mining from Biomedical Literature with Deep Neural Networks*
Supervisors: Dr. Sheng Yu (academic) & Dr. Jiahong Dong (clinical)
Overall GPA Ranking: 2/32

Research Experience

2025– **Research Fellow** (federal employee)
2022–2025 **Visiting Fellow** (intramural trainee)
National Library of Medicine, National Institutes of Health
Supervisor: Dr. Zhiyong Lu

2017–2019 **Visiting Research Scholar**
Department of Biomedical Informatics, University of Pittsburgh
Advisors: Dr. Xinghua Lu & Dr. William W. Cohen (CMU)

Awards & Honors

2025	**Pathway to Independence (K99/R00) Award**	National Institutes of Health
2024	Distinguished Poster Award	AMIA Annual Symposium
2024	Director's Challenge Innovation Award	National Institutes of Health
2023–2024	Fellows' Awards for Research Excellence	National Institutes of Health
2022	Outstanding Graduate Thesis (1/32)	Tsinghua University
2021	Best Clinical NLP Paper	IMIA Yearbook
2020–2021	Top Performance	TREC Biomedical Tracks
2019	First Place Winner	BioBank Disease AI Challenge
2015	National Scholarship	Ministry of Education of China
2015	Gold Medal	iGEM Competition
2013	Gold Medal	China Chemistry Olympiad

Selected Publications

All publications available at Google Scholar (4,500+ citations with 65%+ as the first or co-first author)
† → Equal contribution

Language Model Evaluation

E1. Yang, Y., **Jin, Qiao**, Huang, F. & Lu, Z. Adversarial prompt and fine-tuning attacks threaten medical large language models. *Nature Communications*. https://www.nature.com/articles/s41467-025-64062-1 (2025).

E2. **Jin, Qiao**, Chen, F., Zhou, Y., Xu, Z., Cheung, J. M., Chen, R., Summers, R. M., Rousseau, J. F., Ni, P., Landsman, M. J., Baxter, S. L., Al'Aref, S. J., Li, Y., Chen, A., Brejt, J. A., Chiang, M. F., Peng, Y. & Lu, Z. Hidden Flaws Behind Expert-Level Accuracy of Multimodal GPT-4 Vision in Medicine. *npj Digital Medicine*. https://www.nature.com/articles/s41746-024-01185-7 (2024).

E3. **Jin, Qiao**, Yuan, Z., Xiong, G., Yu, Q., Ying, H., Tan, C., Chen, M., Huang, S., Liu, X. & Yu, S. Biomedical Question Answering: A Survey of Approaches and Challenges. *ACM Computing Surveys*. https://dl.acm.org/doi/abs/10.1145/3490238 (2022).

E4. **Jin, Qiao**, Dhingra, B., Cohen, W. W. & Lu, X. Probing Biomedical Embeddings from Language Models. *Proceedings of the 3rd Workshop on Evaluating Vector Space Representations for NLP (RepEval)*. https://aclanthology.org/W19-2011 (2019).
– Pre-trained BioELMo, the first decoder-only language model in biomedicine.

E5. **Jin, Qiao**, Dhingra, B., Liu, Z., Cohen, W. W. & Lu, X. PubMedQA: A Dataset for Biomedical Research Question Answering. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. https://aclanthology.org/D19-1259 (2019).
– One of the most commonly used benchmarks for evaluating biomedical LLMs.
– Adopted by Google, Microsoft & OpenAI, Meta, and Anthropic.

AI Agents and Tool Learning

A1. **Jin, Qiao**, Yang, Y., Chen, Q. & Lu, Z. AgentMD: Empowering language agents for risk prediction with large-scale clinical tool learning. *Nature Communications*. https://www.nature.com/articles/s41467-025-64430-x (2025).

A2. Tang, X., **Jin, Qiao**, Zhu, K., Yuan, T., Zhang, Y., Zhou, W., Qu, M., Zhao, Y., Tang, J., Zhang, Z., Cohan, A., Greenbaum, D., Lu, Z. & Gerstein, M. Risks of AI scientists: prioritizing safeguarding over autonomy. *Nature Communications*. https://www.nature.com/articles/s41467-025-63913-1 (2025).

A3. Wang†, Z., **Jin†, Qiao**, Wei, C.-H., Tian, S., Lai, P.-T., Zhu, Q., Day, C.-P., Ross, C. & Lu, Z. GeneAgent: self-verification language agent for gene-set analysis using domain databases. *Nature Methods*. https://www.nature.com/articles/s41592-025-02748-6 (2025).

A4. **Jin, Qiao**, Yang, Y., Chen, Q. & Lu, Z. GeneGPT: Augmenting Large Language Models with Domain Tools for Improved Access to Biomedical Information. *Bioinformatics* (2024).

A5. Khandekar†, N., **Jin†, Qiao**, Xiong†, G., Dunn, S., Applebaum, S. S., Anwar, Z., Sarfo-Gyamfi, M., Safranek, C. W., Anwar, A. A., Zhang, A., Gilson, A., Singer, M. B., Dave, A., Taylor, A., Zhang, A., Chen, Q. & Lu, Z. MedCalc-Bench: Evaluating Large Language Models for Medical Calculations. *NeurIPS Datasets and Benchmarks Track*. https://neurips.cc/virtual/2024/poster/97666 (2024).

Literature Search and Retrieval-Augmented Generation

R1. Wang, Z., Cao, L., **Jin, Qiao**, Chan, J., Wan, N., Afzali, B., Cho, H.-J., Choi, C.-I., Emamverdi, M., Gill, M. K., Kim, S.-H., Li, Y., Liu, Y., Luo, Y., Ong, H., Rousseau, J. F., Sheikh, I., Wei, J. J., Xu, Z., Zallek, C. M., Kim, K., Peng, Y., Lu, Z. & Sun, J. A foundation model for human-AI collaboration in medical literature mining. *Nature Communications*. https://www.nature.com/articles/s41467-025-62058-5 (2025).

R2. **Jin, Qiao**, Leaman, R. & Lu, Z. PubMed and Beyond: Biomedical Literature Search in the Age of Artificial Intelligence. *eBioMedicine*. https://www.thelancet.com/journals/ebiom/article/PIIS2352-3964(24)00023-9/fulltext (2024).

R3. Xiong†, G., **Jin†, Qiao**, Lu, Z. & Zhang, A. Benchmarking Retrieval-Augmented Generation for Medicine. *Findings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*. https://aclanthology.org/2024.findings-acl.372/ (2024).
– MedRAG is a widely used toolkit for performing RAG in medicine.
– Regularly maintained and updated, e.g. i-MedRAG and RAG-Gym.

R4. **Jin, Qiao**, Kim, W., Chen, Q., Comeau, D. C., Yeganova, L., Wilbur, J. & Lu, Z. MedCPT: Contrastive Pre-trained Transformers with Large-scale PubMed Search Logs for Zero-shot Biomedical Information Retrieval. *Bioinformatics* (2023).
– MedCPT has been downloaded over 4 million times on Hugging Face.

R5. Zhao†, Z., **Jin†, Qiao**, Chen, F., Peng, T. & Yu, S. A Large-scale Dataset of Patient Summaries for Retrieval-based Clinical Decision Support Systems. *Scientific Data*. https://www.nature.com/articles/s41597-023-02814-8 (2023).

AI for Clinical Trials

C1. **Jin, Qiao**, Wang, Z., Floudas, C., Chen, F., Gong, C., Bracken-Clarke, D., Xue, E., Yang, Y., Sun, J. & Lu, Z. Matching Patients to Clinical Trials with Large Language Models. *Nature Communications*. https://www.nature.com/articles/s41467-024-53081-z (2024).
– TrialGPT is the first framework that utilizes LLMs for clinical trial matching.
– Supported by the NIH Director's Challenge Innovation Award.
– Selected by the AI and Machine Learning focus of *Nature Communications*.
– Featured in the Health Science Top 25 of 2024 in *Nature Communications*.

C2. **Jin, Qiao**, Tan, C., Chen, M., Liu, X. & Huang, S. Predicting Clinical Trial Results by Implicit Evidence Integration. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. https://aclanthology.org/2020.emnlp-main.114/ (2020).
– Selected as the Best Clinical NLP Paper in 2020 by IMIA Yearbook.

Media Coverage

2025	NIH News	NIH researchers develop AI agent that improves accuracy of gene set analysis by leveraging expert-curated databases [A3]
2025	NIH Catalyst	Homegrown AI – NIH Researchers Are Creating AI Tools Benefiting the Broader Biomedical Community [A3, C1]

2024	POLITICO	Using AI to match patients to clinical trials [C1]
2024	NIH News	NIH-developed AI algorithm matches potential volunteers to clinical trials [C1]
2024	AUA News	Connecting Patients to Clinical Trials With Artificial Intelligence [C1]
2024	MedScape	AI's Limitations in Clinical Decision-Making [E2]
2024	NIH News	NIH findings shed light on risks and benefits of integrating AI into medical decision-making [E2]
2024	COSMOS	As we grapple with sovereign AI, perhaps we should treat computational resources as finite and precious [A4]

Academic Service

2025–	Area Chair	ACL Rolling Review (ARR)
2024–	Associate Editor	Journal of Medical Internet Research (JMIR)
2024–	Editorial Committee	Journal of Biomedical Informatics (JBI)
2024	Editorial Committee	Journal of the American Medical Informatics Association (JAMIA) Special Issue

Reviewer for Nature Medicine, Nature Methods, Nature Biomedical Engineering, Nature Computational Science, Nature Communications, Nature Reviews Bioengineering, npj Digital Medicine, Cell Reports Medicine, NeurIPS, ICLR, ICML, SIGIR, ACM MM, ISMB, etc.

Last updated: November 5, 2025