Won Seok Jang

WonSeok_Jang@student.uml.edu | memy85.github.io | Google Scholar | github.com/memy85

PhD candidate building and evaluating LLM conversational agents for patient care: SFT and GRPO post-training of multi-turn agents, simulation-based and human-validated evaluation, and deployment toward a VA clinical pilot. Clinical nursing (BSN) and biomedical-informatics training.

EDUCATION

University of Massachusetts Lowell — Lowell, MA
Ph.D., Computer Science (Miner School of Computer & Information Sciences) — *Sept. 2023 – Expected May 2028*
Advisor: Hong Yu, Ph.D.

Yonsei University — Seoul, South Korea
M.S., Biomedical Systems Informatics, College of Medicine — *Mar. 2021 – Aug. 2023*
Advisor: Yu Rang Park, Ph.D.

Yonsei University — Wonju, South Korea
B.S., Nursing, Wonju College of Medicine — *Mar. 2013 – Feb. 2017*

RESEARCH EXPERIENCE

UMass BioNLP Group, University of Massachusetts Lowell — Lowell, MA
Graduate Research Assistant — *Sept. 2023 – Present*

- **NoteAid patient-educator agent (RL post-training)**: Trained a Qwen3-8B agent that teaches patients their own discharge notes over multi-turn dialogue, using SFT followed by GRPO with a teach-back-based reward. Built AMIE-style automated filtering of the SFT/RL corpus and retrieval grounding over public drug and disease references. *verl*, *vLLM*, *PyTorch*, *Transformers*, *SLURM*.
- **Clinical deployment**: Built the full-stack application that serves the fine-tuned model to patients for a Veterans Affairs pilot: streaming chat, span-grounded highlighting of the source note, session persistence, and researcher dashboard; model runs locally so patient data stays on site. *FastAPI*, *React/TypeScript*, *SQLite*.
- **DischargeBench (evaluation)**: Designed a pre-deployment benchmark in which LLM educators teach persona-grounded simulated patients; curated 320 MIMIC-IV cases across 16 ICD chapters; benchmarked eight LLMs with four metrics, including a claim-level safety audit, scored by an LLM judge validated against physician annotations. The best model reached only 0.70 topic coverage and 0.72 simulated comprehension, and anxious or distrustful patients lowered comprehension for every model.
- **Hallucination in multi-agent systems (evaluation)**: Measured rationale-level hallucination of single- and multi-agent LLM systems across models and clinical tasks against a factuality/faithfulness taxonomy, with human-annotated judge validation. Collaboration improved task accuracy but not hallucination: of 12 resolvable model-by-system comparisons, 8 increased hallucination and 4 decreased it, and a self-correction experiment traced the direction to the underlying model.

Digital Healthcare Lab, Yonsei University College of Medicine — Seoul, South Korea
Research Assistant — *Feb. 2021 – Feb. 2023*

- **Synthetic clinical data**: Built deep generative and Bayesian-network models to synthesize cancer patient records from Yonsei Severance Hospital data under local differential privacy; showed that synthetic data improve survival-status prediction models for early-onset colorectal cancer (*JCO Clinical Cancer Informatics*, co-first author). *PyTorch*, *scikit-learn*, *R*.
- **Federated and distributed learning**: Developed weighted Cox proportional hazards models trained across two hospitals without sharing patient data; contributed to PPFL, a personalized progressive federated learning method for institution-specific features (*iScience*); prototyped cloud-edge split learning on vertically partitioned data. M.S. thesis: score-based voting aggregation for federated neural architecture search. *PyTorch*, *R*, *OpenMined*.
- **Survival analysis (national pharmacovigilance platform)**: Modeled early immunotherapy response in non-small cell lung cancer with time-varying Cox models; found NLR, CRP, and LDH trajectories predictive of 1-year survival (*Translational Lung Cancer Research*, co-first author). *R*, *SQL*.

INDUSTRY EXPERIENCE

Curai Health — Remote
AI Research Intern — *June 2025 – Aug. 2025*

- **Primary-care LLM assistant**: Contributed to LLM assistants that take notes and ask intake questions for patients visiting their primary-care clinician, using supervised fine-tuning.

TECHNICAL SKILLS

LLM training and serving: PyTorch, Hugging Face Transformers, verl (SFT, GRPO), vLLM, retrieval-augmented generation, SLURM/HPC multi-GPU training

Evaluation: Benchmark design, simulated-patient evaluation, LLM-as-a-judge with human/physician validation, inter-annotator agreement, statistical analysis

Engineering: Python, R, SQL, FastAPI, React/TypeScript, SQLite, Git, Linux

Clinical data: MIMIC-IV, EHR notes, VA health data; IRB protocols and data use agreements

PUBLICATIONS

* equal contribution

- **Jang WS**, Jang Y, Ko C, Kim H, Yao Z, Yu H. DischargeBench: Pre-Deployment Evaluation of LLM Discharge Educators through Persona-Grounded Patient Simulation. *Under review, ACL Rolling Review*; October 2026. arXiv:2609.20827
- **Jang WS**, Joshi HB, Yu H. Does Multi-Agent Collaboration Reduce Hallucination in Clinical Tasks? *Under review, ACL Rolling Review*; October 2026. arXiv:2610.27027
- **Jang WS**, Sultana S, Yao Z, Tran H, Yang Z, Kwon S, Yu H. Enhancing Large Language Models for Identifying and Prioritizing Important Medical Jargons From Electronic Health Record Notes Using Data Augmentation: Comparative Study. *JMIR AI*. 2026 Jul 17;5:e75561. doi:10.2196/75561
- Ouyang F, Hu W, Reisman J, Pogoda TK, Carlson KF, Liu W, Moradi Y, Zhang Y, Han S, Sultana S, **Jang WS**, Yao Z, Mitra A, Yu H. Association between PTSD and health-related social needs in US Veterans: an NLP analysis using Veterans Health Administration Data. *Journal of Affective Disorders*. 2026 Feb 15;395:120643. doi:10.1016/j.jad.2025.120643
- **Jang WS**, Tran H, Mistry M, Gandluri S, Zhang Y, Sultana S, Kwon S, Zhang Y, Yao Z, Yu H. Chatbot To Help Patients Understand Their Health. In: *Findings of the Association for Computational Linguistics: EMNLP 2025*; 2025. arXiv:2509.05818
- Yao Z, Sun M, **Jang WS**, Kwon S, Kwon S, Yu H. DischargeSim: A Simulation Benchmark for Educational Doctor-Patient Communication at Discharge. In: *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*; 2025. arXiv:2509.07188
- Tran H, Yao Z, **Jang WS**, Sultana S, Chang A, Zhang Y, Yu H. MedReadCtrl: Personalizing medical text generation with readability-controlled instruction learning. *arXiv*; 2025. doi:10.48550/arXiv.2507.07419
- Yao Z, Parashar A, Zhou H, **Jang WS**, Ouyang F, Yang Z, Yu H. MCQG-SRefine: Multiple Choice Question Generation and Evaluation with Iterative Self-Critique, Correction, and Comparison Feedback. In: *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL), Volume 1: Long Papers*; 2025. p. 10728–10777. doi:10.18653/v1/2025.naacl-long.538
- Yang Z, Yao Z, Tasmin M, Vashisht P, **Jang WS**, Ouyang F, Wang B, McManus D, Berlowitz D, Yu H. Unveiling GPT-4V's hidden challenges behind high accuracy on USMLE questions: Observational Study. *J Med Internet Res*. 2025 Feb 7;27:e65146. doi:10.2196/65146
- Kim TH, Yu JY, **Jang WS**, Heo SC, Sung M, Hong J, Chung K, Park YR. PPFL: A personalized progressive federated learning method for leveraging different healthcare institution-specific features. *iScience*. 2024 Oct 18;27(10). doi:10.1016/j.isci.2024.110943
- Yao Z, Zhang Z, Tang C, Bian X, Zhao Y, Yang Z, Wang J, Zhou H, **Jang WS**, Ouyang F, et al. MedQA-CS: Benchmarking large language models clinical skills using an AI-SCE framework. *arXiv preprint* arXiv:2410.01553; 2024.
- Kim H*, **Jang WS***, Sim WS, Kim HS, Choi JE, Baek ES, Park YR, Shin SJ. Synthetic Data Improve Survival Status Prediction Models in Early-Onset Colorectal Cancer. *JCO Clin Cancer Inform*. 2024 Jan;8:e2300201. PMID:38271642
- Sung M*, **Jang WS***, Kim HR, Park JA, Lim SM, Kim HR, Cho BC, Park YR, Hong MH. Prognostic value of baseline and early treatment response of neutrophil-lymphocyte ratio, C-reactive protein, and lactate dehydrogenase in non-small cell lung cancer patients undergoing immunotherapy. *Transl Lung Cancer Res*. 2023 Jul 31;12(7):1506–1516. PMID:37577328

HONORS AND AWARDS

Best Thesis Award, Korean Society of Medical Informatics (KOSMI) Fall Conference: November 2022

PROFESSIONAL SERVICE

Reviewer, ACL Rolling Review (ARR): May 2025, May 2026, August 2026