

Shane Waxler

Madison, WI • shane@shanewaxler.com • linkedin.com/in/shanewaxler • 858-444-7997

Summary

ML research and engineering and team lead at Epic Systems with 4 years of experience pretraining clinical foundation models and building clinical decision-support tools used by clinicians across 500+ health systems. Lead author of *Generative Medical Event Models Improve with Scale*, detailing Curiosity: a generative medical-event foundation model trained on 118M patients / 151B tokens, establishing medical event model scaling laws across 78 tasks. Research done in collaboration with Microsoft Research and Yale University.

Papers & Media

Generative Medical Event Models Improve with Scale (Curiosity)

LEAD AUTHOR

S. Waxler, P. Blazek, D. White, D. Sneider, ..., R. Shah[†]

arXiv:2508.12104 (2025) - Epic + Microsoft Research + Yale

 [arXiv](#)

- Pretrained 1B-parameter medical event transformer from scratch.
- Developed 78 healthcare evaluations (diagnosis, prognosis, treatment planning, operations).
- Measured pretraining and test-time compute scaling laws for medical event models.

Curiosity Deep Dive: Generative Modeling for Improving Patient Outcomes

INVITED TALK

Stanford BMIR Colloquium

 [Link](#)

Guest Lecture - BMDS 272: Building Successful Startups in Healthcare

Stanford Medicine

 [Link](#)

Experience

Epic Systems — Verona, WI

ML Research and Engineering Team Lead, Curiosity team

Oct 2024–Present:

- Led Curiosity's development and managed a 7-engineer research team to advance model capabilities.
- Led pretraining and post-training for Curiosity models up to 8B parameters.
- Built a platform for prospective validation of clinical models across hundreds of health systems.

ML Researcher, Cosmos

Sep 2022–Oct 2024:

- Pretrained 100M-parameter contrastive patient models for clinical similarity search.
- Fine-tuned LLMs for generalized clinical prediction from longitudinal EHR data.
- Delivered clinical decision-support features to 2,000+ hospitals across the United States.
- Built secure data processing and tokenization pipeline to process 300M+ patient records in <1 hour.

Skills

- **Programming Languages:** Python, SQL, Bash, TypeScript, LaTeX.
- **Libraries:** PyTorch, vLLM, HuggingFace, AzureML, Pandas, Dask.
- **ML/DL:** Pretraining, FSDP, post-training, mixed precision, tokenizers, distributed training.
- **Healthcare:** EHR data, health systems, data privacy & compliance, safe deployment.

Education

Bachelor of Science in Computer Science, Minor in Mathematics

2018–2022

San Francisco State University

Teaching Assistant, CSC 665: Artificial Intelligence, Prof. Pooyan Fazli

Spring & Fall 2022

Student Researcher, People and Robots Laboratory (PeRL)

Studied user perceptions of 73 robot hands through year-long online and in-person human-robot interaction studies.