

Jack Ryan

ryanjack@alumni.stanford.edu

therealjackryan.com

Experience

Member of Technical Staff, Zep AI (YC W24)

Oct. 2024–Present

- Ship production AI software affecting tens of millions of end users.
- Employee #4; contributed to Zep's 40× ARR growth through engineering and enterprise deployments.
- Serve as technical lead for major enterprise deployments, including multibillion-dollar companies.
- Core contributor to Graphiti (30K+ GitHub stars).

Research Assistant, Stanford Trustworthy AI Group (Prof. Fei-Fei Li)

2021

- Co-authored *On the Opportunities and Risks of Foundation Models* (7,600+ citations).
- Researched techniques for neural-network interpretability.

Research Fellow, Stanford Existential Risks Initiative (mentored by Chelsea Voss, OpenAI)

2020–2021

- Built tools to investigate multimodal feature circuits in CLIP.
- Identified three novel neural-network patterns through interpretability experiments.

Education

Stanford University, M.S. Computer Science (AI track)

2023–2024

Stanford University, B.S. Mathematics

2017–2024

Technical Skills

Python, Go, C/C++, SQL, React, Spark/PySpark, PyTorch

Selected Projects

Codebase SCM

Built a versioned structural causal model for software codebases, enabling coding agents to analyze change impact, select affected tests, and simulate interventions from evidence-backed dependency graphs.

[GitHub](#)

[Python](#)

[Causal Modeling](#)

[Developer Tooling](#)

[Agent Systems](#)

CLIP Enrichment Circuits

Discovered 3 novel neural net patterns. Mentored by Chelsea Voss (OpenAI).

[URL](#)

[Github](#)

[Python](#)

[PyTorch](#)

[Captum](#)

[Computer Vision](#)

[Interpretability](#)

Graph Neural Network Fraud Detection

Trained an edge-based graph neural network for financial fraud detection. CS224W final project (2023).

[Colab](#)

[Blog Post](#)

[Python](#)

[PyTorch](#)

[NumPy](#)

[scikit-learn](#)

[Graph Neural Networks](#)

Publications

- [1] Rasmussen, P., Paliychuk, P., Beauvais, T., **Ryan, J.**, and D. Chalef. “Zep: A Temporal Knowledge Graph Architecture for Agent Memory”. In: Preprint in *ArXiv* (2025). 440+ citations. [URL ↗](#)
- [2] Bommasani, [...] **Ryan, J.**, [...] P. Liang. “On the Opportunities and Risks of Foundation Models”. In: Preprint in *ArXiv* (2021). 7,600+ citations. [URL ↗](#)

Honors

AIME (American Invitational Mathematics Examination) Qualifier	2017
2 x Catriona M. Mitchell Mathematics Award	2016, 2017
Hampshire College Summer Studies in Mathematics — Ky and Yu-Fen Fan Scholarship	2016

Testimonial

“Jack drove a complex ML project through leading the design/implementation end to end. Jack is motivated, smart, easy to work with, and detail-oriented. Jack is awesome!”

- *Min Kim, Senior Software Engineer @ Google Labs*

Reference Letter from OpenAI Technical Staff

“I mentored Jack and two other Stanford students on a circuits-style interpretability project using CLIP in 2021–2022.

For historical context, at that time CLIP was a model at the field’s absolute cutting edge when it came to image understanding. Chris, the rest of the OpenAI Interpretability team, and I had been working on some of our own investigations into multimodal neurons in CLIP, so I had context with which to evaluate and mentor this work.

I interacted with Jack and the rest of the project team through regular calls on the group’s progress. Through those calls, I was impressed at the progress the team made, including Jack, who was very motivated and consistently asked insightful, thought-provoking questions during our meetings.

At that time I was thinking a lot about research taste and how to cultivate it in myself, and in the group’s questions and especially Jack’s, I observed a lot of this same virtue: noticing things that we and the field are still confused by, being willing to dig into what ideas for experiments might resolve those confusions, extracting research directions from that, and overall just using that sort of instinct and intuition to pursue scientific hypotheses and generate new science.

Jack made a few novel discoveries about how different units within CLIP contribute to one another, which I found both interesting and valuable. I was also impressed with the technical difficulty of the project, since it involved working with model internals and understanding CLIP’s architecture well. This required the team to be willing to get their hands dirty with running experiments on real model weights.

From my experience mentoring him, I can attest that Jack is curious, hard-working, intelligent, and easy to work with, and he would be a great addition to an AI research team.”

- *Chelsea Voss, Member of Technical Staff @ OpenAI*