

Robert G. Gambee

Portfolio: <https://rgambee.github.io>
robertgambee@gmail.com • (914) 672-3352 • Boston, MA

AI Research Engineer dedicated to improving the world for future generations

Core Competencies

- Proven ability to rapidly master new technologies and adapt in response to evolving needs
- Expertise at driving complex projects from inception to production
- Talent for balancing big-picture strategic thinking and detail-oriented execution
- Experience designing systems to be reliable, scalable and maintainable
- Passion for leveraging technical skills to create meaningful positive impact

Software Skills

Proficient

- Python
- JavaScript, TypeScript, React
- SQL, PostgreSQL, BigQuery
- Git, Jira
- Claude Code, Cursor

Experienced

- PyTorch, scikit-learn
- Docker
- GitHub Actions
- Langfuse, OpenTelemetry
- Bash

Familiar

- Google Cloud Platform
- Amazon Web Services
- Next.js
- Go
- Rust

Professional Experience

FutureSearch: Startup using AI for research and forecasting 2025 to Present
AI Research Engineer

Responsibilities

- Evaluate and enhance capabilities of AI research agents
- Build robust pipelines for running thousands of AI agents concurrently
- Work throughout the stack to rapidly iterate on product ideas
- Log data and build dashboards to monitor performance, stability and cost
- Keep up to date with latest developments in AI capabilities and tools

Technical Projects

- Tech Lead for tools to let AI research agents search the internet and read documents 2025 to Present
 - Developed and evaluated tools to answer queries about web pages
 - Extensively analyzed agent transcripts to identify failure modes
 - Eliminated leading failure mode by switching to superior page retrieval service
 - Optimized page reading flow to reduce maximum latency by factor of 5
- Tech Lead for feature which can rank any list using natural language criteria 2025 to Present
 - Equipped AI agents with tools for enriching the input data to provide an accurate ranking
 - Created evaluations to measure performance on challenging real-world problems
 - Optimized parameters to reduce cost by factor of 10 and run time by factor of 2
- Web app for orchestrating thousands of AI agents to research hard questions at scale 2025 to Present
 - Owned features for importing, viewing and exporting tabular data
 - Designed and built UI for inspecting every step an agent took to produce its answer
 - Prioritized transparency and visibility to help users understand and trust the results
- AI validator for uncertain claims 2025
 - Developed workflow to search the internet for relevant information
 - Created judge to weigh evidence and estimate the probability that the claim is true
 - Wrote and maintained a suite of challenging evaluations
 - Benchmarked custom validator against other approaches across many models

MIT FutureTech – AI Risk Initiative

2026 to Present

Software Engineer (Part-Time Contractor)

- Collaborated with researchers to build databases of AI risks, incidents and mitigations
- Unified and productionized AI pipelines for collecting, screening and classifying documents
- Created visualizations and tools to help decision makers understand and respond to AI risks

Formlabs: Industry leader in professional 3D printing

2015 to 2025

Systems Integration Engineer III

Responsibilities

- Architected software which was maintainable, scalable and testable
- Owned key printer systems for the entire product cycle, driving them from inception to public release
- Understood complex interactions between printer systems, keeping both details and big picture in mind
- Optimized for printer reliability through robust design and failure mode prediction
- Analyzed and visualized printer data to answer pressing questions and inform business decisions
- Mentored junior team members to foster their technical abilities

Technical Projects

- Senior embedded developer for Form 4 and Form 4L 2021 to 2025
 - Thoughtfully crafted powerful yet understandable API to control all aspects of product functionality
 - Comprehensively audited API for security vulnerabilities and coordinated plan to address them
 - Advised architectural decisions for embedded and desktop software
- Data management on user-replaceable components for Form 4 and Form 4L 2023 to 2024
 - Designed a unified architecture for all components, agnostic to interface and data format
 - Wrote extensive validation checks to be robust to failures when reading or writing
 - Secured system against counterfeiting to protect company's primary revenue stream
 - Thoroughly tested all code with automated checks
- Print preparation routine for Form 4 and Form 4L 2021 to 2024
 - Sped up routine by a factor of 5 to 10 compared to previous product, vastly improving user experience
 - Wrote predictive checks to give user advance warning of issues and avoid interrupting prints
 - Implemented specification for how to handle over 50 possible errors
- Prototype firmware for early iterations of Form 4 2021 to 2022
 - Independently developed prototype firmware in Python to support crucial conceptual testing
 - Balanced competing desires for flexibility and stability using a modular design
 - Rapidly responded to feature requests and bug reports, addressing them in days if not hours
- Dashboard for plotting live sensor data 2021
 - Independently developed over four days during company hackathon
 - Wrote backend in Go, wrote frontend in JavaScript, streamed data via WebSockets
 - Recognized by the CEO in a company-wide email as one of the most impressive projects that year

Achievements

- Recipient of Formlabs' Perform Award, which recognizes top 10% of employees 2020 and 2023

Independent AI Research

Do LLMs have stable preferences?

2025 to 2026

- Built upon findings in Claude 4 System card to understand how preferences vary under different conditions
- Used pairwise comparisons between tasks to determine the LLM's implicit rating for each task
- Compared multiple models, response formats, orders and numbers of tasks
- Found that ratings for individual tasks combine coherently when considering sequences of tasks
- Uncovered strong dependency on response format: free-form vs. structured output
- Results have implications for AI welfare, highlighting challenges with satisfying LLM preferences

- Reproduction of “The Capacity for Moral Self-Correction in Large Language Models”** 2023
- Loaded and processed tens of thousands of samples from three different datasets
 - Submitted API requests asynchronously, with automatic retries and rate limiting
 - Analyzed bias in model responses according to three different metrics
 - Compared and contrasted results to demonstrate influence of RLHF training vs. prompt engineering

Volunteer Experience

- AI Governance and Safety Canada** 2024
- Designed and implemented flexible system for scraping information relevant to the AI safety community
 - Leveraged LLM to robustly scrape many sites with zero site-specific tuning required
 - Initially developed to find upcoming events, but can be extended to other types of content, e.g. publications
 - Created automated workflow using GitHub Actions to run scraper and publish output to database
 - Overhauled collection of introductory AI resources with updated list covering many topics and formats

Personal Projects

- Chronicle** 2023
- Web app to keep track of how one spends one’s time*
- Used Django framework to manage HTTP requests and access SQLite database
 - Presented data as a table for sorting and filtering, as well multiple charts for visualization
 - Set up automated test and deployment workflows using GitHub Actions

- SCAFFOLD** 2023
- Completed as part of AI Safety Camp (3 person team plus advisor)*
- Built React web app to generate feedback on one’s research ideas using GPT
 - Fine tuned model to make its responses more relevant to AI safety research

Publications

FutureSearch: Nikos I. Bosse, Jon Evans, **Robert G. Gambee**, Daniel Hnyk, Peter Mühlbacher, Lawrence Phillips, Dan Schwarz, Jack Wildman; Deep Research Bench: Evaluating AI Web Research Agents. 6 May 2025. <https://doi.org/10.48550/arXiv.2506.06287>

Jonas L. Kaufman, Scott H. Tan, Kirklann Lau, Ashka Shah, **Robert G. Gambee**, Chris Gage, Lupe MacIntosh, Albert Dato, Peter N. Saeta, Richard C. Haskell, Todd C. Monson; Permittivity effects of particle agglomeration in ferroelectric ceramic-epoxy composites using finite element modeling. AIP Advances 1 December 2018; 8 (12): 125020. <https://doi.org/10.1063/1.5053442>

Education

- Harvey Mudd College**, Claremont, CA
- Bachelor of Science in Engineering with High Distinction* 2011 to 2015
- GPA: 3.8
 - Inducted into Tau Beta Pi, national engineering honor society 2014
 - Recognized on Dean’s List of top performing students 2012 to 2015