

# Devin Fitzpatrick

[devinfitzpatrick@berkeley.edu](mailto:devinfitzpatrick@berkeley.edu) | [linkedin.com/in/devinfitzpat](https://www.linkedin.com/in/devinfitzpat) | [github.com/Devin-Fitzpatrick06](https://github.com/Devin-Fitzpatrick06)

## EDUCATION

### University of California, Berkeley

GPA: 3.68

*B.A. Computer Science + B.S. Business Administration*

*Graduation: May 2028*

- Courses: Data Structures & Algorithms, Linear Algebra, Techniques of Data Science, Structures & Interps. of CS, Optimization Models in Engineering, Discrete Mathematics & Probability Theory, Computer Architecture, Efficient Algorithms & Intractable Problems (IP), Probability & Random Processes (IP), Intro to Machine Learning (IP)

## EXPERIENCE

### CLEAR

June 2026 – August 2026

*Software Engineering Intern*

*New York City, NY*

- Reduced PR review time by 60% by deploying an AI-powered GitHub App on Kubernetes (EKS) invoking AWS Bedrock/Claude via TypeScript to automate inline code review comments across 5+ engineering teams
- Secured LLM infrastructure for 5+ teams by provisioning AWS Secrets Manager and IAM roles/policies on EKS
- Drove 3x AI tooling adoption by building a PR Review Handler enabling 10+ teams to innersource LLM workflows

### Berkeley Artificial Intelligence Research (BAIR)

August 2026 – Present

*Undergraduate Researcher, Darrell Group*

*Berkeley, CA*

- Isolated 3 root causes of VLM hallucination by benchmarking LLaVA and Qwen2.5-VL on 4 datasets in PyTorch
- Quantified a 22% CHAIR increase under visual degradation by scripting blur, occlusion, and resolution ablations
- Attributed 30% of hallucinations to language priors across 500+ false-premise prompts on MMHal and HaloQuest
- Cut evaluation runtime 40% by parallelizing VLM inference on 8 A100 GPUs via vLLM, SLURM, HuggingFace

### Launchpad

September 2025 – Present

*Machine Learning Engineer*

*Berkeley, CA*

- Engineered intelligent software for real-world problems by collaborating on 3+ student and industry projects
- Developed autonomous systems using Computer Vision and Reinforcement Learning, achieving 85%+ accuracy
- Implemented 10+ generative AI and language models to optimize user experiences by prototyping 3+ models

### TurbineOne

June 2025 – August 2025

*Software Engineering Intern*

*San Francisco, CA*

- Benchmarked 15+ AI startup datasets by synthesizing 6+ months of LinkedIn data to identify 5+ market trends
- Improved internal asset retrieval speed by 30% for 100+ assets by redesigning indexing with metadata schemas
- Automated outreach and engagement pipelines to reduce manual data entry by 60% and scale signal analysis

## PROJECTS & OPEN SOURCE

### INFRA-MAXXER | *TypeScript, Next.js, React, Recharts, REST APIs, Vercel*

July 2026

- Cut projected GPU spend 16.8% by building a TypeScript optimizer pricing usage against 180-day forward curves
- Engineered 5 REST API routes in Next.js fusing Ornn forward quotes with 120 days of Datadog cost telemetry
- Shipped a React sandbox re-pricing 6 GPU types on 4 live inputs, backtested over 12 months of trailing spend

### Whisk (Matcha Making Robot) | *Python, PyTorch, Gemini, OpenCV, MuJoCo, LeRobot, NumPy*

May 2026

- Built Gemini LLM agent pipeline in Python orchestrating 5+ real-time tool calls to control an SO-101 robot
- Engineered pose estimation pipeline using OpenCV solvePnP and AprilTag detection, cutting grasp error by 30%
- Implemented inverse kinematics and collision-aware path planning in simulation, achieving 90%+ task completion

### Oculy | *Python, PyTorch, PANDAS, SciPy, NumPy, OpenSignals, PyGame, Matplotlib*

December 2025

- Achieved 85%+ eye-direction classification accuracy by training a 1D CNN on noisy EEG data in Python/PyTorch
- Built low-latency deep learning inference pipelines with sub-50ms response times for real-time signal processing
- Integrated neural network gaze controls into 3 real-time prototypes, enabling hands-free HCI across 2+ use cases

## SKILLS

**Programming:** Python, TypeScript, JavaScript, Java, SQL, C++, React, Next.js, Flask, Tailwind CSS, PyTorch, NumPy, pandas, scikit-learn, OpenCV, MuJoCo, AWS (Bedrock, EKS, IAM), Docker, Kubernetes, GitHub Actions, Git

**Other:** Claude Code, Codex, Vercel, Jira, Figma, Microsoft (Excel, PowerPoint)

**Interests:** Running, Spikeball, Basketball, SciFi Movies (e.g. Marvel), Music, FIFA, Lifting Weights, Snowboarding