

# James Maa

E-mail: jamesmaa17@gmail.com • Phone: (650) 388-0628 • Website: jamesmaa.com • Current Location: San Francisco, CA  
LinkedIn: linkedin.com/in/jamesmaa • Github: github.com/jamesmaa • Goodreads: goodreads.com/jamesmaa  
Spoken languages: English (Native), Mandarin Chinese 中文 (Fluent), Japanese 日本語 (Proficient, JLPT N1), Spanish (Proficient)

## WORK EXPERIENCE

---

### Career Break

Jan 2024 - Present

- Grew my open-source project, Yomitan (popup dictionary for language learning) from 10,000 to 90,000+ users within a year
  - Added support for multiple languages, including whole word scanning, language-specific settings and grammatical logic
  - Led a team of collaborators to prioritize removing user friction for non-technical users by adding onboarding and UX features
- Personal development and travel

### Discord

Aug 2020 - Jan 2024

#### Senior Software Engineer, Platform & Services

Remote + Hybrid @ San Francisco, CA

- Shipped multiple projects to improve the reliability, velocity, and usability of the deployment system tied to our monolithic API server that 200+ engineers interact with every day and service x billions of requests a day. These projects result in going from 2-3 deploys (each taking multiple hours) a week performed by a small group of trusted engineers to 10+ continuous deployments a day, increasing throughput by 2000%, reducing deploy times by 90%, and shrinking the average changeset size by 80%
  - Engineered and built an automatic db migration tooling and integrated it across all of our environments (dev, CI, staging, production), providing a smooth local dev UX working with migrations in Postgres and Cassandra and automating hundreds of dangerous and manual db operations in production, removing significant operational burden from the database team
  - Transitioned our monolithic API server from running on virtual machines in Salt to containers in kubernetes
  - Designed and rolled out a new deployment strategy onto Buildkite to automate the CI/CD process of the API servers, enabling developers to move safer and faster
  - Introduced deployment analysis that evaluates deployment health of deployments based on custom metrics, removing the need for a trained human operator to shepherd the deployment and saving engineers 30+ hours a week
- Migrated Discord's legacy asynchronous task system in Celery to GCP Pub/Sub by providing a easy-to-user developer interface for defining asynchronous tasks and automating the managing, scaling, and monitoring of the underlying infrastructure resources
- Drove team and company wide reliability efforts through new telemetry instrumentation, monitoring, dashboards, and processes.
- Owned all the API-related infrastructure and platforms, including our Kubernetes clusters, Terraform, Cloudflare workers, Celery, Pub/Sub, Redis, and Google Cloud Platform, serving as the catchall team for everything infrastructure and DevOps at Discord.

### Scratch Services (Fintech/Loan Servicing)

April 2018 - Dec 2019

#### Software Engineer, Backend & Infrastructure

In person @ San Francisco, CA

- Implemented, tested, and released concurrency in our core accounting infrastructure to increase the system concurrency level from 1 to 30 and expand our core system capacity by more than 1800%.
- Delivered multiple mission-critical loan mechanic features like principal disbursements and forbearances on top of a custom-built temporal database with a transaction-log-inspired service handling out-of-order delivery of order-dependent computation, managing \$X00,000,000 of debt
- Automated multiple manual business operation tasks that saved the company more than 30 hours a week of manual work including automated onboarding, loan sales, and adjustments.

### Reddit

April 2016 - April 2018

#### Software Engineer, Ads

In person @ San Francisco, CA

- Wrote the campaign management backend from scratch that is the source of truth for tens of thousands of ad campaigns resulting in a 350% increase in revenue from self-service advertisers
- Remodeled the payments and delivery data models to stomp out rampant data discrepancy issues in our original payment system. The new billing service successfully onboards hundreds of new advertisers and processes \$X00k/week of charges
- Designed and wrote the inventory forecasting service that uses time series machine learning models to predict inventory and is used for pacing for all of Reddit's X,000's of ad campaigns

### Internships

Summers of 2013 - 2015

- 2013 @ Yahoo!: Wrote a data processing ETL to generate novel statistical insights for tens of thousands of online merchants
- 2014 @ Pinterest: Designed, trained, and productionized an adaptive logistical regression machine learning model that detects suspicious logins, improving recall by 16x and caught 42x more suspicious logins than the previous implementations
- 2015 @ Pinterest: Product Management Intern managing a team of 6 engineers shipping Pinterest's SDK to public release

## EDUCATION & SKILLS

---

### UC Berkeley

### Skills

B.A. Computer Science  
Aug 2012 - Dec 2015  
GPA: 3.7

Mastery: Python, Kubernetes, Google Cloud Platform, Cassandra, CI/CD, Terraform, Buildkite, Pub/Sub  
Proficient: Networking & traffic, SRE, Datadog, Observability  
Cold storage proficient: Ruby, Chef, AWS, Postres, Jenkins, React, Node, Javascript