

PRANAV DHINGRA

New York, NY | +1 (224) 714-9724 | pranav.dhingra@gmail.com | [LinkedIn](#) | [Portfolio](#)

6 years across product management and engineering. Delivered \$49.5M in annual savings optimizing ML systems at Grubhub. Cornell Tech MEng CS, focused on LLMs, Deep Learning, and Reinforcement Learning. Founder of No-Code Works, teaching 90+ students to build with AI.

EXPERIENCE

Senior Product Manager / Product Manager II

2022 – 2025

Grubhub — San Francisco, CA / Chicago, IL

- Launched Priority Delivery: designed exposure experiments, wrote backend requirements for dispatch prioritization. 550K+ orders, 46% adoption, \$12M ARR. Maintained late delivery targets (<10% priority, <20% standard) while optimizing network-wide fulfillment cost.
- Drove \$49.5M annual savings by coordinating ETA policies with auto-dispatch optimization algorithm. Designed multi-variate experiments across three teams, navigated trade-offs between fulfillment cost, user experience, and care costs. Increased bundled orders by 10%.
- Optimized how ML predictions are presented to users: ran A/B tests on ETA range formats to balance conversion vs. delivery reliability. Widening ranges by 5 minutes reduced late deliveries by 9.28% with no conversion impact. Key insight: users prefer reliable predictions over precise ones.
- Led ML initiative to improve food prep time predictions. Identified ground truth data quality issues through error analysis, proposed new target variable using driver departure as proxy signal, wrote PRD with manual tuning controls, designed A/B and switchback experiments.

Software Engineer II / Software Engineer

2019 – 2022

Grubhub — Chicago, IL

- Built experimentation microservice from scratch: 400K requests/minute, <10ms latency. Implemented in-memory caching with fanout refresh, architected async integration for latency-sensitive services. Primary code-owner; shaped service roadmap by gathering cross-functional requirements and authoring RFCs.
- Scoped A/B testing functionality for diner experiments: gathered cross-functional requirements, authored RFC with implementation timeline and technical specification.

PROJECTS

- [Efficient Attention in Vision Transformers](#) — Comparative study of Linformer, Performer, and Nyströmformer attention in ViTs. Implemented full architecture from scratch in PyTorch; proposed CNN+Linformer hybrid that recovered 5-13pp accuracy while reducing compute 17-31%. Evaluated on ImageNette (10-class ImageNet subset).
- [No-Code Works](#) — Founded AI builder community at Cornell Tech (90+ members). Lead sessions on Claude Code, Cursor, and Vercel. Teaching technical and non-technical students to ship products using LLM-powered development tools.
- [Village](#) — Childcare coordination platform for parents to form trusted communities. Product-engineered cross-platform app (iOS, Android, Web) in 4 weeks using Claude Code, Bolt, Supabase. Instrumented with PostHog to track engagement and guide iteration.

EDUCATION

Cornell Tech, Cornell University — M.Eng. Computer Science, New York, NY

2025 – 2026

Coursework: Language Modeling, Deep Learning, Generative Models, Reinforcement Learning, Trustworthy AI

Northwestern University — B.A. Computer Science & Economics, Evanston, IL

2015 – 2019

SKILLS

AI/ML: LLMs, Deep Learning, Transformers, Attention Mechanisms, PyTorch, Prompt Engineering, Claude Code, Cursor

Engineering: Python, SQL, AWS, Microservices, Distributed Systems

Product: Experiment Design, A/B Testing, Switchback Testing, Cross-functional Leadership