

Vinay Shashidhar

Director of Machine Learning · AI-Native Products Leader

Bengaluru, India · Open to global relocation (visa-ready) · [linkedin.com/in/vinayshashidhar](https://www.linkedin.com/in/vinayshashidhar)

EXECUTIVE SUMMARY

Machine Learning leader with 15+ years building and scaling AI functions from 0→1 across FinTech, gaming, travel, and e-commerce. Three times stood up an ML organization from scratch — most recently driving a company-wide shift to AI-native engineering at ClearTax, shipping 16 products and revenue-generating AI infrastructure in a single year. Inventor on a granted US patent and published at ACL, KDD, and ICML; equally comfortable setting AI strategy for the board and shipping production code.

WORK EXPERIENCE

ClearTax — Director of Machine Learning

Sep 2024 – Present

India's largest B2B tax & compliance SaaS platform — 4,000+ enterprises, 6M+ taxpayers.

- Built ClearTax's AI/ML function from the ground up with a company-wide mandate to make Clear AI-native — leading 10+ AI initiatives and shipping **16 products in FY26** across customer AI, revenue operations, e-invoicing, and internal adoption.
- **Conversational Intelligence** (replaced Gong) — scaled support-call auditing from ≈25 calls/week reviewed manually to **1,200–2,000 calls/day at 100% coverage (≈350x)** across 4 verticals, with real-time fatal-call alerts and a leadership dashboard; ≈\$140K/yr cost avoidance.
- **AI Support Chatbot** — B2C/B2B assistant handling **225K+ conversations at up to 92% deflection** (≈140 FTE-equivalent); cut query-to-ticket from 27% to under 10% and lifted CSAT 3.7 → 4.1.
- **SalesHub** — revenue-intelligence platform (account overview, AI pitch generation, automated post-call evaluation) live for 15+ enterprise clients; >80% adoption; cut AI pitch-generation cost from \$0.50 to \$0.14 per pitch.
- **Global E-Invoice AI** — agent auto-generating positive/negative/edge test cases pre-development, cutting compliance-to-dev from 7–10 to 2–3 days and new-country go-live from ≈1 month to under 1 week (Belgium, Poland, Croatia, Germany, UAE).
- **Agentic Performance Review** — multi-agent system for AI-augmented reviews, live across ≈800 employees; projected ≈50% reduction in review effort.
- **Company-wide AI adoption** — 100% of engineers on agentic coding tools, 80%+ of code agent-generated, 50+ models/agents in production; authored the "AI-First" strategy and three-pillar framework adopted in the board deck.
- *Environment:* LLMs / RAG / agent orchestration (Claude Agent SDK), ASR (Whisper → Gemini), Python, Node/TypeScript, PostgreSQL, Redis, MongoDB, AWS SageMaker.

Hike Labs — Director of Machine Learning

Mar 2022 – Sep 2024

ML leader for Rush — a high-concurrency, skill-based competitive gaming platform with millions of users.

- Established the ML division from scratch (engineers, scientists, analysts) for a real-time, high-concurrency consumer platform.
- Built real-time ML using generative modeling for **fair matchmaking & skill assessment**, increasing new-user (N1) retention; skill framework shared with national gaming regulators.
- Engineered predictive churn for highest-LTV users (**drop-off –4%**) plus personalized next-best-action; built a real-time anomaly-detection / early-warning platform for integrity and revenue.
- Spearheaded an emotionally intelligent **Hindi voice-to-voice support bot** (speech-to-text, text-to-speech, speech-to-speech).
- Team impact: >2% **absolute retention**, **8–10% ARPU lift**, **25% lower op-cost** → **2–3% absolute CM1%**, at 99.99%+ reliability.

Rakuten Institute of Technology — Staff Data Scientist

Sep 2020 – Mar 2022

Corporate research lab of the global e-commerce and internet-services group.

- Led the Audience Science team within customer marketing (behavioral analysis, predictive modeling).
- Built look-alike modeling with fastText / sent2vec embeddings; newsletter recommendations via LightFM; a predictive lead-scoring algorithm.
- Proposed & led the **AI Research Lab** initiative — weekly deep technical discussions on the latest AI/ML research across multiple teams.

OYO Rooms – Machine Learning Engineer III

Aug 2018 – Sep 2020

Global hospitality marketplace operating across 35+ countries.

- First member of the ML team; established the pipeline and a data-driven engineering culture.
- Built a **dynamic-pricing** algorithm (revenue maximization as a quadratic program) → 4–5% net-revenue increase.
- Built demand + cancellation models (overbooking / auto-prepaid) → 3–5% net-revenue increase by cutting cancellations 10%; built a surge-pricing engine.
- Estimated price elasticity with University of Malaya researchers; modeled customer payment-method preference to maximize recovered revenue.

IBM Research – Software Engineer · India Research Lab

Jun 2016 – Jun 2018

- Delivered the feedback module of a chat-based **Cognitive Tutor** (response classification + dialogue strategy) in a major pilot with a leading education firm; quantified effects of tutor misclassification on student behavior (published at ITS & EDM).
- Applied ML to neuroscience (Power-Up): analyzed fMRI + behavioral data to study emotion's effect on brain state.

Aspiring Minds – Senior R&D Engineer

Jun 2010 – Jan 2016

- Co-developed an automated **Spoken English evaluation** product (speech recognition + ML) that became market leader at 70–80% share.
- Managed teams building automatic evaluation of computer programs (static + dynamic analysis) and a first-of-its-kind automated psychomotor-skill assessment.
- Built automated regional skill-mapping / job-skill-gap analysis that opened new-country accounts; built the first personality inventory in the Indian context and a PSO-based on-job-performance classifier — a company differentiator.

PATENTS

US9984585B2 — Constructed response grading **GRANTED**

WO2015022637A1 — Candidate response evaluation for job fitment

US20190362642A1 — Tutoring agent improvement recommendation

EDUCATION

Executive Masters, ML & Data Science

IIT Hyderabad · 2017–2019

B.E., Computer Science

NSIT, University of Delhi · 2006–2010

SELECTED PUBLICATIONS

Impact of Tutor Errors on Student Engagement in a Dialog-Based ITS — **ITS 2018**

Automatic spontaneous speech grading: a novel crowd feature-derivation technique — **ACL 2015**

Spoken English grading: ML with crowd intelligence — **KDD 2015**

Learning models for personalized educational feedback & selection — **ML-ED, ICML 2015**

Principles for using ML in assessment of open-response items — **NeurIPS workshop**