

SHAIMAK REDDY

Senior AI Product Manager | Agentic & Voice AI · Product Strategy

+91-7799-70-5555 | shaimak.reddy@gmail.com | linkedin.com/in/shaimakreddy | shaimakreddy.com

SUMMARY

Senior AI Product Manager with 9 years of experience, 7 building Voice and Agentic AI at Samsung Bixby (400M+ Galaxy devices). Co-led Samsung's first Agentic Assistant to the Galaxy S26 launch, authored the quality standard Bixby ships against, and now drive agentic platform strategy at the executive level.

IIM Bangalore (Director's Merit List) | IIT Bombay (All India Rank 3)

EXPERIENCE

Senior AI Product Manager | Samsung Electronics, Seoul HQ

Mar 2023 - Present

Own product strategy for agentic task management, AI quality and evaluation, and multilingual global rollout of Bixby / Galaxy AI.

- **Agentic Ecosystem Strategy:** Defined the vision and roadmap to move Bixby from a reactive voice assistant to an autonomous, multi-step Agentic Task Layer that acts across third-party apps and the web.
- **Roadmap & Org Leadership:** Own the roadmap and quarterly reprioritization for the agentic platform and related features through to executive sign-off; mentor 4 PMs across Bangalore and Mountain View as HQ decision point for the US and India teams.

Agentic Platform (Galaxy S26, 2026)

- **Galaxy S26 Agentic Launch:** Lifted complex-task completion 15% by co-leading the launch of Samsung's first Agentic Assistant, executing requests end-to-end across multiple steps.
- **Agentic Reliability:** Cut unresolved requests from 25% to ~6% by owning the recovery layer no single team had: a disambiguation step for low-confidence intents and a Perplexity-backed fallback for out-of-scope queries.
- **"Second Brain" Recall:** Captured 25%+ of previously failing queries (35% adoption in six months) by setting the on-device RAG strategy for personal queries (messages, notes, calendar).

Quality & Architecture

- **LLM Evaluation Standard:** Raised platform quality from 5.3 to 6.7/7 in a quarter and set the company-wide 6.5/7 release bar by authoring the LLM evaluation standard (binary pass/fail, replacing subjective 1-to-5 scoring).
- **Bixby Chat Re-Engagement:** Doubled utterances per active user per week and lifted monthly active users about 50% by rebuilding Bixby Chat with AI-generated responses (replacing templates) across 14 languages (Galaxy S25, 2025).
- **Natural-Language Device Help:** Cut settings failures from 40% to under 10% and deflected 15,000+ support tickets a day by shipping Device QnA, a RAG knowledge base over device manuals and past support tickets.
- **Cross-Region Delivery:** Cut cross-team dependency resolution from over a week to under 36 hours by embedding on-site for a joint sprint across the US, Korea, and India, aligning 80 PRDs.
- **Privacy & Data Governance:** Partner with Legal on every release and own privacy-by-design review of personal-data handling across the US, EU, and Korea.

AI Product Manager | Samsung Electronics, Seoul HQ

Aug 2019 - Feb 2023

Scaled Bixby from 150M to 300M+ users as consumer AI PM.

- **On-Device Bixby (Founding PM):** Delivered 35% server-cost savings, ~30% faster responses, and 90%+ US and Korea adoption by scoping that 20 core commands covered half of all usage and setting the on-device size target from competitive benchmarks.
- **Personal Voice Creator:** Reached 20M users in six months with Samsung's first consumer voice-cloning feature, built fully on-device because voice is sensitive biometric data, setting Samsung's privacy-by-design standard (9to5Google, TechRadar).
- **Text Call:** Scaled a new AI calling feature to 50M+ calls in 6 months across 1B Samsung phones; it became a foundation for Galaxy AI, with a fast post-launch fix lifting NPS 20 points (9to5Google, SamMobile).
- **Bixby Labs:** Founded Bixby's first opt-in platform for experimental AI features, now the org's standard launch-and-learn channel; first release Custom Wake-up drove 17% activation and 25% higher weekly actives vs. control (9to5Google, TechRadar).
- **Hands-Free Engagement:** Put 40% of commands within reach from the lock screen by turning a security team's "no" on voice unlock into a three-tier, risk-based access model; added Continuous Conversation for wake-word-free follow-ups to deepen engagement.

Business Analyst | ElectrifiAI (formerly Opera Solutions), Delhi

Jul 2015 - May 2017

B2B data analytics for enterprise and government clients, including a top-3 global payment network and a Big-4 auditor.

- **Tax-Evasion Risk Modeling (Govt. of India):** Cut on-ground investigation effort 30% by risk-profiling 1.8M flagged accounts from \$175B in demonetization deposits with KNN and rules-based ML to surface likely evaders.
- **High-Risk Targeting:** Enabled \$130M+ in seizures (toward \$1.5B) by narrowing to a 1,000-account shortlist via statistical outlier detection for the government's raids.

EDUCATION

IIM Bangalore, MBA

Jun 2017 - May 2019

Director's Merit List Awardee (top 5% of batch)

Summer Associate, Deutsche Bank: built an interest-rate index POC (1.25+ Sharpe ratio).

IIT Bombay, Bachelor of Technology, Electrical Engineering

Jun 2011 - May 2015

All India Rank 3 in IIT-JEE (450,000+ candidates)

Research Intern, IIM Ahmedabad: Built Python NLP models to achieve 90.8% sentiment analysis accuracy on book reviews.

TECHNICAL SKILLS

AI & ML: Multi-agent orchestration, agentic task planning, LLM evaluation & benchmarking, RAG, LoRA fine-tuning, ASR / NLU / TTS.

Product Frameworks: Hybrid on-device/cloud voice pipelines, prompt engineering, privacy-by-design, experiment & success-metrics design, pass/fail quality evaluation.

Tools & Platforms: Python, Linux, Docker, self-hosted LLMs (Llama), OpenClaw, open-source stack (Outline, Mattermost, Jellyfin).