

Shiv Palit

📍 New York City, NY ✉ shivpalit@gmail.com 🌐 in/shivpalit 🌩 shivpalit.com 📞 912-334-6950

SUMMARY

Product Manager with a strong Data Science foundation and extensive experience in ML and Data Engineering. I lead AI and ML initiatives in fast moving environments by driving product vision, directing cross functional teams, and translating complex problems into clear direction that accelerates execution and delivers meaningful impact.

SKILLS

Data Science & ML: Python, SQL, PyTorch, Scikit-learn, Pandas, LLMs, RAG, Streamlit

Data Engineering: Databricks, Airflow, DBT, Snowflake, Dagster, Spark

Cloud & Infrastructure: AWS, GCP, Heroku, Docker

Product Management: Product strategy, roadmap design, experimentation, KPI development, cross functional leadership

EXPERIENCE

Sr. Product Manager, Warner Bros. Discovery – New York, NY June 2023 – Present

- Led the CX Data & AI team, delivering scalable data products that unified customer insights across support, social, and app ecosystems to inform operational and product strategy.
- Co-led the AI Evaluation Program, defining the product strategy and operating model for assessing internal and third-party AI solutions across CX. Owned roadmaps, UAT, ramp, and go/no-go criteria for production decisions.
- Launched WBD's first AI knowledge assistant for CX Advocates, owning roadmap and rollout. Built the knowledge repository and Google-based retrieval layer, reducing handle time by 10% and saving \$2M annually.
- Delivered \$1.5M in annual savings by owning the end-to-end product design of a real-time LLM classification system, including decision logic, Databricks orchestration, and automated resolution of non-actionable cases.
- Drove a 95% reduction in localization turnaround time by launching a self-service translation platform across 6+ markets and 30+ languages, saving \$50K per regional HBO Max launch.
- Built a standardized ML platform by designing a PyTorch-based training API that automated PII anonymization, Salesforce ingestion, and ensemble model training, enabling faster iteration and repeatable model development.

Founding Data Engineer, Obi9 Technologies – New York, NY Nov 2024 – Oct 2025

- Built the core data engineering infrastructure at Obi9, an early-stage AI-insights platform providing real-time ingestion and automated analysis of biotech-related financial documents for institutional clients.
- Defined the 12-month product and engineering roadmap for the platform's beta, aligning development milestones with market needs and investor expectations.
- Supported fundraising efforts by shaping the business plan, investor narrative, product market fit, and serving as the technical lead alongside the founding team.
- Designed and implemented a structured knowledge repository to track fact-level changes across SEC filings, earnings transcripts, and conference documents, forming the foundation for LLM-driven signal detection.
- Developed a proprietary graph-based extraction engine and all supporting backend APIs to transform unstructured text into interconnected fact entities and insights for downstream processing.
- Architected and deployed the company's AWS cloud infrastructure with Dagster-orchestrated pipelines to ensure scalable, reliable, and secure distributed data processing.

Analytics Consultant, Das42 Consulting – New York, NY Oct 2021 – June 2023

- Led the implementation of a full-stack BI and data infrastructure for a B2C services client using Looker, GCP, and Airflow, enabling a 285% increase in annual subscription revenue (to \$100M+).
- Served as the primary client partner - managing delivery, expectations, and long-term analytics strategy - and drove over \$1M in revenue for Das42 by securing three contract extensions.
- Enabled data-driven acquisition of 130K+ active subscriptions by developing cross-functional data models and Looker dashboards linking acquisition cost, retention, and lifetime value (LTV).
- Developed churn and LTV prediction models using Python and SQL, creating new retention KPIs and branch-level performance metrics through engineered subscription and behavioral features.
- Built scalable data pipelines using GCP and Airflow to unify subscription, marketing, payroll, timecard, and call center data into a centralized data warehouse.

Data Science Engineer, Pura Vida Investments – New York, NY Apr 2020 – Oct 2021

- Built and automated data pipelines using Heroku, PostgreSQL, Redis, and Python to aggregate 25+ alternative data sources, reducing manual research overhead and saving analysts an average of 6 hours per day.
- Engineered interactive research dashboards in Streamlit to centralize datasets, automate analytics workflows, and deliver real-time insights for investment research teams.
- Deployed large-scale NLP pipelines that ingested and analyzed 100K+ daily posts and comments from online forums, performing sentiment modeling to track shifts in public perception of portfolio companies.
- Designed and implemented quantitative clustering systems using PCA, Dynamic Time Warping, and KMeans to identify groups of assets with correlated trading behaviors and uncover hidden market relationships.
- Developed signal detection models using alternative data sources including aviation activity, social media sentiment, and web scraped operational metrics to surface early indicators of portfolio company performance.

EDUCATION

Johns Hopkins University – MS in Data Science 2025
Davidson College – BS in Economics 2019