

BHARATH KUMAR

+1 (201) 790-6958 | bkumar.datascientist@gmail.com | [LinkedIn Profile](#) | North Bergen, NJ (NYC Metro)

Data scientist and AI engineer building production agentic systems and ML pipelines that help Fortune 500 advertisers scale with Amazon Advertising & AWS Clean Rooms. 8+ years across data engineering, ML, and production software.

SKILLS

- **AI / Agentic Engineering:** Claude Code, Kiro, Gemini Antigravity, LangGraph, CrewAI, Strands, MCP/CLI, RAG, context / prompt / loop / memory-layer engineering, LLM & agent evaluation, agent-directed development
- **Advertising Domain:** Amazon Marketing Cloud (AMC), AMC Custom Models, AWS Clean Rooms, Amazon DSP, Social Media Ads Platform (Meta, Snapchat, X, LinkedIn, Reddit, Pinterest, TikTok), Google Ads Platform (SA360, DV360, Google Search Console, GA4), Measurement Platforms (Nielsen, DoubleVerify, LiveRamp), Marketing Science & Attribution.
- **Technical:** Python, SQL, TensorFlow, Scikit-learn, Tableau, RESTful APIs, Docker, Amazon Web Services (Clean Rooms, SageMaker, Bedrock, AgentCore, Step Functions, Lambda, CloudFormation, CloudWatch, CloudFront, API Gateway, DynamoDB, Athena, S3, EC2), Productivity Tools (MS-Office Suite, JIRA, Asana, Quip)
- **Data:** System Design, Data Engineering (ETL, Data Migration, Workflow Automation), Data Visualization and Reporting, Data Analytics, Data Wrangling, Statistical Inference, Data Modeling (Decision Tree, K-Means Clustering, SVM, Gaussian Naïve Bayes, Neural Network), Machine Learning (Supervised and Unsupervised Learning), Deep Learning, NLP, A/B Testing

EXPERIENCE

Amazon Advertising

New York, NY

Ad Tech Consultant / Solutions Architect (MADS Advanced Solutions)

Mar 2022 - Present

- Sole architect of an Agentic Ad-Ops Builder Kit, a production platform (harness + skills + memory + CLI) that orchestrates AMC ML pipelines, DSP campaigns, and measurement through conversational AI, built on AWS CDK, Bedrock AgentCore, Strands, Step Functions, Lambda, and DynamoDB.
- Designed and led AMC Custom Models through closed beta as technical lead for enterprise clients (P&G, Kimberly-Clark, Nespresso), architecting end-to-end onboarding and hand-building the baseline ML code (training/inference Python, SQL, Docker) for new-to-brand, repeat-purchase, and lifetime-value models, producing 1,750+ modeled audiences and 806M ad impressions, with published case-study results (Nespresso: 67% lower CAC, 213% more new-to-brand).
- Architected the onboarding and automation for AMPC, building reusable CLF templates and Lambda/EventBridge automation that onboarded all priority (P0) measurement vendors onto AWS Clean Rooms.
- Built the PAD partner-delivery framework (idea → production model) that 4+ teammates replicated team-wide; flagship deliverable was the CLTV model (cited by [AdExchanger](#); adopted by partners to win net-new subscriptions).
- Led paid-feature validation bridging engineering and customers across multiple closed betas like Audience Segment Insights, Amazon Garage, NCS Offline Insights, Swift Instance Creation, Category Insights, driving features to GA.

Hearts & Science (Omnicom Media Group)

New York, NY

Senior Data Scientist / Associate Director of Data Engineering (Marketing Science)

Jul 2018 - Mar 2022

- Built AT&T's data-engineering operation from scratch, architected a production data warehouse and automated data ingestion across several ad & measurement platforms, cutting processing time ~98% (weeks → ~2 hours) into client-facing dashboards.
- Designed an in-house Python ingestion framework that onboards complex new data feeds in minutes with near-zero maintenance, standardizing engineering/data-science workflows and cutting ~30% of project dev time.
- Co-built a custom bid-optimization model (regression + S-curve over propensity) on the AppNexus/Xandr ad server, with API interfaces for data collection and model deployment.
- Built time-series models to detect creative-behavior patterns and predict landing-page views.
- Prototyped a taxonomy-correction process (from internal-client requirements) that fixed cross-team data inconsistencies, making downstream joins faster and more reliable.

Grapple Media Inc.

New York, NY

Freelance Data Scientist

Aug 2017 – Nov 2018

- Built an NLP misinformation classifier (IBM Watson, Google Cloud, MS Cognitive Services) to filter fake news and surface quality content at scale.
- Engineered the supporting data foundation: a real-time news-scraper API, automated processing pipelines, and an AWS RDS store, handling everything from ingestion to modeling.

EDUCATION

Stevens Institute of Technology

Hoboken, NJ

Master of Science in Computer Engineering | Major: Data Engineering, Data Science

Jan 2018

Awards: “Best Hack for Social Good” - Awarded by *JP Morgan Chase*; “Alexa Skill Developer” - Awarded by *Amazon*

Anna University

Chennai, Tamil Nadu, India

Bachelor of Engineering – Electronics & Communication Engineering

May 2016