

SHUBHAM JOSHI

Data Scientist and Machine Learning Engineer | San Francisco, CA | (202) 440-8445 | shubhamjoshipro.mail@gmail.com

[Tacet, live product](#) | [Live model](#) | [Research report](#) | [Tableau exhibits](#) | [Portfolio](#) | github.com/shubhamjoshipro/mail-svg | linkedin.com/in/shubham-joshi1

SUMMARY

Data scientist who turns unreliable data and ambiguous labels into defensible decisions and deployed models. On a year long client engagement I reconciled four inconsistent federal sources, caught a twelvefold error that nothing had flagged, isolated a policy effect at p below 0.001 with robust standard errors, and presented the implications in person to a national organization's leadership. I fine tuned and published a transformer that went from 52.3 to 73.1 percent by fixing the labels rather than the model, and I ship what I build: 508 commits and 92,448 lines across 11 public repositories, 8 of them live. Now I am building Tacet, my own company, where a learned rule only counts once it holds on cases the expert never showed.

SKILLS

Data Science, Statistics and Experimentation: Data science across healthcare, financial services and consumer settings; econometric analysis, OLS and logistic regression, HC3 robust standard errors, causal impact and attribution, hypothesis testing and significance interpretation, experimental and quasi experimental design, A/B testing, cross validation, ROC-AUC and macro and weighted F1, forecasting, clustering and segmentation, Monte Carlo simulation, mixed integer linear programming, leakage prevention, error analysis, statistical communication; identifying the key concepts and processes that drive an outcome, and understanding where a model stops being trustworthy

Machine Learning and Deep Learning: Python (pandas, NumPy, scikit-learn, statsmodels, PyTorch), Transformers and Hugging Face, ModernBERT fine tuning and public deployment, LightGBM, XGBoost, CatBoost, H2O AutoML, Optuna, SMOTE and class imbalance handling, feature engineering, Platt calibration, PSI drift monitoring, SHAP and LIME interpretability, fairness and equity auditing, cost based threshold tuning, abstention

Data Engineering, Tools and Development: Applied computer science and programming; SQL (PostgreSQL, Oracle; complex joins, CTEs, window functions), multi source reconciliation across inconsistent schemas and geographic grains, ETL and ELT pipeline design, data quality validation and error detection, data dictionaries and documentation, reproducible pipelines and version control with Git, FastAPI, Prisma, Docker, CI/CD, Railway and Hugging Face deployment, large and complex datasets; designing and developing innovative analytical tools and reusable solutions as deployed services

Analysis, Visualization and Communication: Framing an ambiguous business question, translating analysis into business implications, Tableau dashboards and exhibits, charts, written research reports with prespecified evaluation protocols, memos and documentation, executive and client readouts delivered in person, explaining technical findings to non technical audiences, collaborating in small analytical teams while working independently

EXPERIENCE

Founder, Tacet - San Francisco, CA

Oct 2026 - Present

- Founded Tacet and built it solo from a HackNation and ElevenLabs hackathon prototype into a live multi tenant product. It watches an expert work, predicts each decision before they make it, asks one question only when it cannot explain what they did, and compiles the answer into an executable rule in the expert's own words that AI agents can call over API or MCP and new hires are coached against.
- Made the learning provable rather than claimed: a plain Python engine with Bayesian thresholds where language models only translate, sealed tests whose predictions are SHA-256 committed before any human label, and a compile gate that rejects a rule unless it quotes the expert verbatim and reproduces their past decisions. On a sample ledger, one answered question took held back accuracy from 96 to 100 percent across 44 transactions.

Analytics Consultant, Johns Hopkins University, Client: Alliance of Community Health Plans - Washington, DC

Jul 2025 - Jul 2026

- Led most of the analysis for a six person team answering a federal reimbursement question nobody had measured, for a national organization of nonprofit health plans. Built the dataset it required: a four stage Python and SQL pipeline reconciling four federal sources published on different schedules and geographic grains into 138,263 matched plan county records across 208 variables, and caught a twelvefold enrollment overcount caused by mismatched geography before it reached a single model or slide.
- Isolated the causal cost impact with OLS and HC3 robust standard errors controlling for enrollment and quality rating, finding premiums higher by 2.8 dollars and maximum out of pocket higher by 97 dollars, both significant at p below 0.001, plus a five fold stratified logistic model at 0.642 cross validated AUC forecasting market exit and k-means segmentation validated three ways.
- Delivered a competitive benchmark establishing a client cost advantage in 493 of 875 counties, Tableau exhibits, a data dictionary and a written report presented in person to leadership. The client said it would change where the organization spent its resources; the team placed top three school wide.

Co-Founder and Vice President, Business Analytics and AI Organization, Johns Hopkins University - Washington, DC

Jan 2026 - Present

- Co-founded and lead a 100 member organization alongside my degree, delivering a 3 hour hands on analytics and machine learning workshop to 35 registrants with the guide published open source.

Founder, Digital Marketing and Web Studio, self-employed - Nagoya, Japan

Jan 2021 - May 2025

- Ran a studio single handed for four years for international clients in a third language, owning scoping, pricing, delivery and reporting, including audience research and A/B testing on live campaigns.

MODELING PROJECTS

OpsPilot, Fine Tuned Transformer Deployed with Abstention

[Live model](#) | [Code](#)

- Fine tuned and published a ModernBERT ticket router on 23,747 leakage safe examples: baselines put TF-IDF at 44 and ModernBERT at 52.3 percent, so the label space was the ceiling, not the model, and rebuilding ten overlapping labels into seven operational queues lifted validation accuracy to 73.1 percent at weighted F1 0.713. Confidence gating auto routes 62.3 percent of tickets at 85.5 percent accuracy.

Credit Risk, Propensity and Valuation Modeling

[Code](#)

- Modeled borrower default risk for a consumer lending portfolio at 0.6857 out of fold ROC-AUC against a 0.5041 baseline, with Platt calibration so scores read as probabilities, PSI drift monitoring, an equity audit and a decision threshold set on expected cost rather than accuracy. Separately placed 7th of 59 teams competing individually on a financial services propensity model at 0.8285 ROC-AUC and 19th on a property valuation regression at 8.708 MAE, both scored on held out data.

RxCheck, Research Report with a Prespecified Evaluation Protocol

[Live app](#) | [Research report](#)

- Normalized 152,416 drug interaction records with exact, brand, fuzzy and NDC matching, kept detection deterministic with the model restricted to grounded explanation, and wrote a research report with seven prespecified evidence packages that gives the unfavourable findings equal billing, including an official benchmark it passed only 22 of 30.

EDUCATION

Johns Hopkins University, Carey Business School - MS, Business Analytics and Artificial Intelligence, STEM designated, Dean's Scholarship Aug 2026

Nagoya University of Commerce and Business - BBA in Management, GPA 3.63 (WES evaluated), Dean's Scholar honors, Monash exchange 2025

Coursework: Statistical Analysis, Advanced Business Analytics, Practical Machine Learning, Generative AI, Simulation, Data Visualization. **Certifications:** SAS SQL Essentials, SAS AI Foundations, Azure AI-900, Anthropic Claude 101. **Languages:** English, Japanese, Nepali, Hindi.