Kalyan Venkatesh

Data · Operations

adavivenkatesh@gmail.com kalyanvenk.com linkedin.com/in/kalyan-venk github.com/kalyan-venk

SKILLS

Languages & Databases: Python, SQL, PostgreSQL
Data Modeling & Transformation: Data Modeling, dbt, Snowflake, ETL
Data Quality: Data Validation, Reconciliation, Query Optimization
Analysis & Experimentation: pandas, NumPy, Statistical Analysis, Experimental Design
Business Intelligence: Tableau, Data Visualization, Automated Reporting
Development & Delivery: Git, API Integration, Requirements Gathering, AI-Assisted Development
Certifications: AWS Certified Cloud Practitioner

EXPERIENCE

sensen.ai | Software Engineer — Jan 2023 – Aug 2024

- Developed 300+ customer-facing data products for 80+ clients as Sensen's sole data scientist & translated stakeholder requirements into reusable SQL functions & materialized views that powered on-demand reporting without manual data extraction.
- Built a company-standard model evaluation framework for data validation across 26 parking deployments, with matched & unmatched vehicle entry/exit views that exposed incomplete sessions & helped reduce unbillable events by ∼5%, with ∼$6K protected monthly per site.
- Automated invoice preparation in SQL through reports that calculated infringement counts, charges & totals for client-selected dates, sites & cameras, reduced work from ∼10 to ∼2 hours per client per 2-week billing cycle & reached nearly the entire client base.
- Built officer productivity metrics for Brisbane City Council from ∼2,000 weekly incidents & hours worked to automate weighted rankings for ∼60 officers over client-selected periods & replace ∼2 hours of manual reporting each week.

sensen.ai | Software Engineer Intern — Aug 2022 – Jan 2023

- Developed reusable PL/SQL functions for ∼15 ticket types that retrieved infringement details across multiple tables from an event ID for ticket-generation APIs to centralize repeated retrieval logic & cut query execution time by ∼50% from 800–900 ms to 400–450 ms.

PROJECTS

Product Compass: Data Pipelines and Analytics for AI Developer Tools

- Built a Python API ingestion workflow into PostgreSQL for 1,159 Codex issues & verified 100% source-to-database reconciliation across 16,226 field comparisons with no duplicate records.
- Design dbt models & Snowflake history tables for 2 AI-product sources with data quality tests & documentation, expected to prevent duplicate records & maintain 100% source reconciliation after incremental updates & recovery.
- Develop Tableau dashboards across 8 issue categories & spreadsheet scenario analysis for 2 triage approaches to quantify expected review throughput & unresolved backlog under explicit capacity & cost assumptions.

Graduate Research: Statistical Analysis of Language Model Performance

- Performed statistical analysis of 17,000 evaluation records across 10 experimental phases in Python, with paired comparisons & 95% confidence intervals to quantify performance changes & variation across 3 trials.
- Measured a +1.37-percentage-point average gain in first-attempt code correctness across 9 conditions for one model but a decline from 90% to 86% for another, which demonstrated the need for model-specific validation.

EDUCATION

DePaul University — Chicago, IL
Master of Science, Computer Science — Sep 2024 – Jun 2026
Visvesvaraya National Institute of Technology — Nagpur, India
Bachelor of Technology, Engineering — Jul 2017 – May 2021