TIM ZHANG

Data Engineer | AI-Driven Pipeline Architect

weijianlucky@gmail.com | +86 131 6809 0613 / +852 84965467

GitHub: tim-portfolio.github.io/tim-portfolio

PROFESSIONAL SUMMARY

Data Engineer with 8+ years of experience specializing in enterprise data warehouse design, cloud migration, and AI-driven data pipeline development. Led migration of 10TB+ historical data to Tencent Cloud, achieving 87% processing efficiency improvement. Currently focusing on LLM-based intelligent system development with production-ready RAG and Agentic architecture expertise. Proven track record delivering high-impact data solutions for Fortune 500 clients including Walmart and Bank of Ningbo.

TECHNICAL SKILLS

Programming & Scripting: Python, SQL (BigQuery/SparkSQL/Impala), Bash/Shell

Big Data & Analytics: Hadoop, Apache Spark, Apache Hive, StarRocks, Kafka

Cloud & DevOps: Google Cloud Platform (Cloud Run, Cloud SQL), Tencent Cloud, Docker, Linux

AI/ML Engineering: FastAPI, LangChain, pgvector, OpenAI API, RAG Architecture, GPT-4o

Data Modeling: Dimensional Modeling, 5-Layer EDW Architecture, DBT

KEY PROJECTS

SRR Agentic Case Processing System | *Technical Lead (99.4% contribution)* | *Lingnan Cup Competition Project* | *2025-2026*

GitHub: github.com/SRR-Project-Team/SRR-Agentic-Case-Processing-System

Problem: Hong Kong Architectural Services Department(ArchSD) case processing relied on manual entry and routing, with average processing cycles lasting days and high error rates.

Solution:

- Designed 7-layer Agentic processing architecture with 17 pluggable atomic capabilities
- Built multi-format document parsing pipeline with field extraction (A-Q categories) for ICC 1823, TMO, RCC channels
- Integrated OpenAI GPT-4o with text-embedding-3-small for intelligent classification and response generation
- Implemented pgvector-based semantic case retrieval with RRF fusion ranking

- Developed 3-tier quality assessment system (L1 keyword → L2 rules → L3 RAGAS LLM-as-Judge)
- Created self-healing mechanism with Best-of-N selection and 4-category failure differential rollback
- Deployed on Google Cloud Run with Cloud SQL and Docker for auto-scaling

Results: 67 commits, 80,430 lines of code across two iterations | 12/12 requirements fully implemented | End-to-end automation achieved | Reduced processing time by 75% (40 min → 10 min) | Projected annual cost savings of HK$6.39M

Tech Stack: Python 3.11+, FastAPI, React + TypeScript, PostgreSQL 15+ pgvector, OpenAI GPT-4o, Google Cloud Run + Cloud SQL, Docker

RAGagent - Lingnan University Policy Q&A System | *Developer* | *2026*

GitHub: github.com/terzotian/RAGagent

Problem: University staff and students struggled to find accurate answers to policy and course-related questions from scattered documentation.

Solution:

- Built semantic search pipeline using pgvector embeddings + TF-IDF with Reciprocal Rank Fusion (RRF)
- Integrated Google Vertex AI (Gemini) with text-embedding-004 for generative responses
- Implemented streaming output for real-time answer generation
- Developed user management, file upload, and conversation history features

Results: 4 GitHub stars, 37 commits, 86.3% Python codebase | Production-ready deployment | Policy Q&A accuracy >95% | Average retrieval time reduced by >70%

Tech Stack: React + TypeScript + Vite, FastAPI, PostgreSQL + pgvector, Google Vertex AI (Gemini)

PROFESSIONAL EXPERIENCE

AI Engineer / System Architect (Part-time) | Lingnan University — SRR & RAGagent Projects | 09/2025 - Present

• Designed and built a 7-layer Agentic architecture with 17 pluggable capabilities for HK ArchSD case processing, delivering 67 commits and 80,430 lines of code with end-to-end automation

• Built RAGagent university policy Q&A system with pgvector + Vertex AI (Gemini), achieving >95% policy Q&A accuracy

• Deployed containerized AI services on GCP Cloud Run with PostgreSQL + pgvector for semantic retrieval

Data Engineer | Shenzhen RenruiHR Technology Co., Ltd. | 09/2023 - 09/2025

• Migrated IDC data warehouse to Tencent Cloud by restructuring into 5-layer ODS/DIM/DWD/DWS/ADS architecture, optimizing data flow and management

• Leveraged Seatunnel for incremental/full data syncs, ensuring 99.95% data consistency with business systems and reducing sync failures by 70%

• Automated workflow scheduling with DolphinScheduler, maintaining 99.7% on-time job execution and freeing 5+ hours/week of engineering time

• Optimized storage with ORC format, slashing disk space consumption by 38% and boosting query speed by 55% for 80% of business-critical tables

• Implemented partitioning/bucketing strategies across non-ADS layers, reducing full-table scan occurrences by 90% and lowering average query resource usage by 45%

• Designed star-schema dimensional models to automate time-based aggregations, cutting manual summarization efforts by 60%

• Enhanced ETL performance with SparkSQL, reducing data processing time from hours to minutes (87.5% improvement)

• Built Python-driven data flow monitoring, reducing quality and timeliness issues by 60%

• Integrated StarRocks with Hive external tables, achieving sub-second query performance (average 0.8s), up from 3-10 minutes previously

Data Engineer | Hangzhou Yatop Information Technology Co., Ltd. | 04/2021 - 08/2023

• Architected scalable 5-layer offline data warehouse (SA/SHM/PDM/BDM/DIM) on Hive, streamlining data processing pipelines and increasing cross-team data reusability by 50%

• Led data integration from 100+ transaction systems using Informatica PowerCenter

• Developed Impala SQL/HQL logic using window functions for metric calculations, accelerating report generation by 40%

• Rationalized table structures and standardized naming, cutting data issue-tracing time by 45%

• Designed daily/monthly partitioned data flows, boosting statistical analysis efficiency by 30%

• Created custom UDFs/UDTFs for complex computations and sensitive data masking, ensuring 100% compliance with data privacy policies

• Enabled auto-merge of small files, decreasing file count by 65% and improving cluster efficiency by 30%

Data Engineer | Shenzhen Owned Technology Co., Ltd. | 02/2019 - 03/2021

• Developed and maintained data model layer of Bank of Ningbo's data warehouse, supplying data for downstream financial reporting

• Processed data with SQL as requirements from data market

• Liaised directly with clients for issue resolution, improving communication skills

Engineer | Chengdu HQtimes Technology Co., Ltd. | 07/2017 - 11/2018

- Troubleshot network problems and system errors for IBM software applications
- Developed Ops scripts with BASH SHELL/Python
- Wrote technical operation documentation and reports

EDUCATION

Lingnan University, Hong Kong | 2025 - 2026

MSc in Artificial Intelligence and Business Analytics

Shenzhen University, China | 2014 - 2017

Bachelor of Management

LANGUAGES

- Mandarin (Native)
- Cantonese (Native)
- English (Professional Working Proficiency)