

REYNALD ACE PILPIL

AI / ML Engineer · Agentic LLM Systems & Document AI

0968 619 1930 | azepilpil139@gmail.com | Cavite, Philippines

GitHub: github.com/aces-14 | Hugging Face: huggingface.co/acepilpil

EDUCATION

Technological University of the Philippines – Cavite

Bachelor of Engineering Technology, Major in Computer Engineering

TECHNICAL SKILLS

Machine Learning & AI: Python, PyTorch, Hugging Face Transformers (BERT, DistilBERT, LayoutLMv3, Flan-T5), scikit-learn, XGBoost, SHAP, sentence-transformers, ChromaDB, NumPy, pandas

LLM & Agentic Systems: LangChain, LangGraph, RAG architecture, multi-agent orchestration, Groq (Llama 3.3 70B), Tavily Search API, Pydantic v2, prompt engineering, LangSmith

MLOps & Deployment: FastAPI, Docker, Hugging Face Spaces, Gradio, Railway, REST API design, Swagger / OpenAPI, pdfplumber, python-docx

Backend & Data: Django, Flask, MySQL, SQL Query Optimization, JavaScript (ES6+)

Tools & Practices: Git/GitHub, Version Control, Agile, Code Review, Testing, Debugging

EXPERIENCE

Technical Sales Intern · 3E HITECH SOLUTIONS, INC.

January 2026 – May 2026

- Supported pre-sales activities for the company's solutions, focusing on projects with the National Grid Corporation of the Philippines (NGCP) — contributing to proposal preparation, scoping, and technical clarifications with stakeholders
- Assisted project management across NGCP deployments — coordinating timelines, tracking deliverables, and supporting communication between technical teams and clients

KEY AI / MACHINE LEARNING PROJECTS

Job Application Assistant Agent *Autonomous Multi-Step LLM Agent*

Tech: Python, Groq (Llama 3.3 70B), Tavily Search API, Pydantic v2, Gradio, python-docx ·

Live: [HF Space](#) · [GitHub](#)

- Built a 9-step agent with **integrity enforcement** — architecturally prohibited from fabricating resume content; a post-tailoring diff check flags violations, with every change logged in a fully auditable changelog (reason + confidence score)
- Pipeline: resume parsing → JD scoring → company research via live web search → ATS-tailored resume + cover letter + LinkedIn outreach delivered as downloadable Word documents in under 30 seconds

Market Intelligence Agent *Stateful Multi-Agent System with Self-Revision*

Tech: Python, LangGraph, LangChain, Groq (Llama 3.3 70B), Tavily Search API, Pydantic v2, Gradio, LangSmith ·

Live: [HF Space](#) · [GitHub](#)

- Built an autonomous research agent with a **self-evaluating revision loop** using LangGraph — agent critiques its own draft, scores it, and re-searches targeted gaps until quality threshold is met, demonstrating stateful multi-node orchestration beyond simple LLM chaining
- Researches a company across four parallel angles (overview, news, competitors, market trends) via live web search and synthesizes the results into a structured intelligence report delivered as a Word document

Philippine Labor Law RAG Assistant *End-to-End Retrieval-Augmented Generation*

Tech: Python, LangChain, ChromaDB, Hugging Face Transformers, Flan-T5, sentence-transformers (all-MiniLM-L6-v2), Gradio ·

Live: [HF Space](#) · [Colab](#)

- Designed and deployed an end-to-end RAG pipeline over the Philippine Labor Code — document ingestion, semantic chunking, vector embeddings (all-MiniLM-L6-v2), ChromaDB retrieval, and Flan-T5 generation
- Returns natural-language answers (leave entitlements, overtime rules, termination procedures) with source citations grounded in the actual law, reducing hallucination through retrieval-first design

Financial Document Intelligence API *Production REST API*

Tech: Python, FastAPI, pdfplumber, pandas, Docker, Railway ·

Live: [HF Space](#) · [GitHub](#) · [API Docs](#)

- Production-grade REST API that ingests bank statement PDFs and returns structured JSON — extracted transactions, auto-categorized spending, computed financial summaries (income, expenses, savings rate), anomaly flags, and actionable insights
- Containerized with Docker and deployed to Railway with auto-generated Swagger / OpenAPI docs, demonstrating end-to-end fintech-relevant deployment

Receipt Data Extractor (LayoutLMv3) *Multimodal Document AI*

Tech: Python, Hugging Face Transformers, LayoutLMv3, PyTorch ·

Live: [HF Space](#)

- Extracts structured data (menu items, quantities, prices, totals) from receipt images using LayoutLMv3, a multimodal transformer that jointly reasons over text, layout, and visual features
- Achieved **93.96% F1 on validation set**, demonstrating document AI capability beyond plain-text NLP

BERT Sentiment Analyzer *Social Media Intelligence Tool*

Tech: Python, Hugging Face Transformers, PyTorch, DistilBERT, Gradio ·

Live: [HF Space](#) · [Colab](#)

- Multi-input bulk sentiment analysis tool that processes batches of reviews, tweets, or comments and returns an aggregate sentiment score, breakdown chart, and the top positive / negative phrases driving the score
- Designed dashboard-style output (not just single-label classification), making it usable as a social-media intelligence tool

Customer Churn Predictor *Customer Retention Dashboard*

Tech: Python, scikit-learn, XGBoost, SHAP, pandas, Gradio ·

Live: [HF Space](#) · [Colab](#)

- Production-ready churn risk dashboard: ingests CSV customer data, predicts churn probability with XGBoost, and outputs a retention report with risk distribution, ranked high-risk customers, and SHAP-based explanations of feature contributions
- Translates raw model output into business-actionable retention recommendations using **explainable AI (SHAP)**