

Pradyumn Tendulkar

Boston, MA | pradyumnten.work@gmail.com | [7745190324](https://www.linkedin.com/in/p-tendulkar) | <https://www.linkedin.com/in/p-tendulkar> | <https://github.com/pradyten>

Education

Worcester Polytechnic Institute, Worcester, MA

Graduation Date: Dec 2025

MS in Data Science (GPA: 4.0)

Courses: Deep Learning, Natural Language Processing, Statistical Methods for Data Science, MLOps, Business Intelligence

Vellore Institute of Technology, India

Aug 2021 - May 2024

BTech in Computer Science and Engineering (GPA: 3.56)

Courses: Python, Machine Learning, Natural Language Processing, Software Engineering, Database Management Systems

Skills

LLM & GenAI: OpenAI API, Claude API, Gemini, Mistral, LangChain, LangGraph, Agentic AI, RAG, Pinecone, Fine-Tuning, Prompt Engineering

ML & NLP: PyTorch, TensorFlow, Keras, Hugging Face, Sentence-Transformers, BERTScore, Scikit-learn, XGBoost, spaCy, NLTK, Gensim

Cloud & MLOps: AWS (Bedrock, SageMaker, Lambda, Step Functions, Textract, Comprehend, CloudFormation), GCP, Vercel, W&B, Docker, Kubernetes, CI/CD

Languages: Python, SQL, TypeScript, R, C/C++

Data & Tools: Pandas, NumPy, Pydantic, FastAPI, Streamlit, Tavily, Playwright, Pillow, PostgreSQL, MongoDB, Git

Experience

AI Engineer, CrossingLegal, Boston, MA

Aug 2025 – Dec 2025

- Architected multi-agent validation pipeline using **AWS Step Functions, Lambda, and Bedrock Claude API** for vision-based PDF extraction into structured JSON, reducing H-1B processing from 5 days to 10 minutes.
- Constructed RAG fallback system with **Titan embeddings**, in-memory vector store, and **cosine similarity retrieval** to resolve ambiguities across 50+ Form I-129 fields, achieving 95-99% extraction accuracy.
- Integrated LLM validation layer using **OpenAI, Pydantic** schema enforcement, and **Textract**-based KB construction, cutting manual verification effort by 70-90%.

AI Engineer Intern, Boredm LLC, Phoenix, AZ

Jun 2025 – Aug 2025

- Fine-tuned **GPT-4o-mini** via **OpenAI API** on 1000+ balanced examples for 58-entity NER extraction, implementing **Pydantic** schema validation to achieve 99.1% production accuracy on soil lithology descriptions.
- Engineered synthetic data pipeline using **Python-TypeScript** hybrid architecture with dependency-aware generation, keyboard-proximity corruption, and 12 structural variants, reducing manual annotation time by 90%.

Research Assistant, Worcester Polytechnic Institute, Worcester, MA ([Publication: Link](#))

Oct 2024 – May 2025

- Developed zero-shot summarization framework using **GPT-4, Gensim GloVe embeddings**, and **NLTK** preprocessing to generate abstractive summaries from 5000+ Amazon reviews, published at **AMCIS 2025**.
- Attained 0.9417 cosine similarity with Amazon's proprietary summaries using **scikit-learn** cosine similarity, **iterative prompt refinement**, and **BERTScore** validation on **Playwright**-scraped review data.

AI Engineer Intern, SmartStream Technologies – India

May 2023 – Jul 2023

- Deployed end-to-end NER pipeline using **AWS Textract, SageMaker Ground Truth** (custom annotation template), and **Comprehend Custom Entity Recognition** to extract entities from native PDF/Word documents, achieving 99% accuracy.
- Replaced legacy rule-based extraction with **Comprehend ML workflows** and **CloudFormation**-provisioned annotation infrastructure, reducing entity extraction cycles from 4 months to 3 weeks (80% improvement).

Projects

RAG-Powered Legal Research Engine

[Live Demo](#)

- Created full-stack RAG application using **LangGraph** (6-node pipeline), **FastAPI, GPT-4o**, and **Pinecone** for legal case research, with dual-layer confidence scoring (60% retrieval + 40% LLM self-assessment) and **Mistral** fallback, hosted on **Vercel/Railway**.

AI Content Pipeline with Self-Improving Hooks

[Live Demo](#)

- Designed multi-agent pipeline using **Claude, Gemini**, and **Tavily** APIs with 3-phase writing chain (research → draft → refine) and self-improving hook analysis that learns from engagement metrics, aggregating 14 news sources into ready-to-post content.

Generalizable Vision-to-JSON Extraction Framework

[Live Demo](#)

- Built vision-based PDF extraction system using **OpenAI Vision API** (GPT-4.1), **pypdfium2**, and **Pillow** with adaptive DPI rendering, supporting 12 document types (I-129, passports, resumes) with JSON schema enforcement, served on **Hugging Face Spaces**.

Predictive Modeling of Water Turbidity Using Remote Sensing Indices

- Co-authored paper (accepted at **ICFAiSE 2025**) achieving $R^2=0.98$ using **LSTM** (**TensorFlow/Keras, PyTorch**) for turbidity prediction through 2-year time-series analysis of Indian lakes, outperforming ARIMA and linear regression baselines.