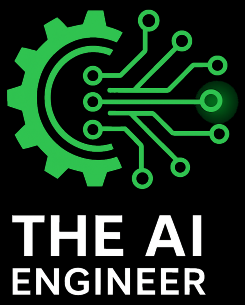

The AI Engineer

A New Standard for Learning and Applying AI

October 6, 2025

Contents

Why Now

AI systems already perform a large share of entry-level work. Support is handled by chatbots; first drafts of legal and marketing text are produced by LLMs; code assistants generate boilerplate and tests; data copilots summarize spreadsheets and dashboards. The result: fewer traditional junior tasks, and more leverage for people who can design, build, and operate AI systems.

1 Executive Overview

The AI Engineer is a modern learning and execution program focused on building end-to-end AI capabilities—combining software engineering, machine learning, LLMs, evaluation, and deployment. It is concise, rigorous, and practical, enabling you to ship real systems quickly while developing strong fundamentals.

High demand and strong compensation. Organizations in every sector are hiring for AI engineering roles—product companies, consultancies, finance, healthcare, and the public sector. Roles commonly advertise six-figure base salaries in major markets, reflecting the scarcity of hands-on AI engineering talent.

2 How AI Changed Work

Entry-level tasks have been automated or reshaped:

- *Support and operations:* chatbots resolve routine tickets; human agents focus on edge cases.
- *Writing and analysis:* LLMs draft emails, briefs, and reports; humans edit and decide.
- *Software:* code copilots propose implementations and tests; engineers review and integrate.
- *Data work:* AI copilots query, summarize, and visualize; analysts validate and drive action.

This shift raises the bar: junior roles shrink, while **AI engineering**—the ability to build reliable AI products—becomes the No. 1 skill for modern teams.

Market dynamic. As standard entry-level tasks are automated away, fewer traditional junior roles remain. At the same time, demand for hands-on AI engineering grows faster than supply. The result is a persistent supply–demand gap and rising compensation for capable AI engineers.

3 The Talent War

This supply–demand gap is visible in compensation and recruiting:

- *Wall Street Journal:* recruiters offering $300M packages over four years to top researchers; source.
- *The New York Times:* young AI researchers recruited like NBA stars with nine-figure packages; source.
- *Wall Street Journal:* tech's hiring spree for AI talent as a once-in-a-generation opportunity; source.

4 Learning Paths

Choose the schedule that fits you; all paths converge to the same skill set.

Fast Track (4 weeks) Full-time, intensive. Live orientation sessions (recorded). Rapid skill acquisition—ideal for a bootcamp pace.

Medium Track (8 weeks) Part-time friendly. Roughly two weeks per module for more practice and deeper exploration.

Self-Paced (6 months) Learn at your own speed with full content access; optional grading and maximum flexibility.

5 Dual-Track Curriculum

Master both theoretical foundations and practical engineering through parallel tracks.

5.1 Core Track

- **Python & Math for ML/DL** *(vectorization, linear algebra, calculus, optimization)*—hands-on Python applied to ML/DL tasks.
- **Deep Learning with PyTorch**—datasets, models, training loops, regularization, schedulers, and evaluation.
- **LLM from Scratch**—tokenization, attention, architecture, training, and sampling to demystify transformers.
- **AI Agents & Automation**—tool use, planning, orchestration, and production cost/latency control.

5.2 Engineering Track

- **Foundation Primers**—Python best practices, math refreshers, and development environment setup.
- **Software Engineering**—testing, packaging, CLI tools, CI/CD pipelines, and containerization.
- **ML Engineering**—data management, experiment tracking, metrics, evaluations, and reproducibility.
- **AI Engineering**—model serving, latency optimization, observability, safety, and privacy in production.

6 A Unique Learning Approach

The program is organized around **seven proprietary, coherently crafted books** that interlock into a practical pathway. Each chapter includes exercises; each book features test projects; and the track culminates in capstones that integrate the full stack (from data and prompting to evaluation, orchestration, and deployment).

What You Will Learn

Foundations: modern software patterns for AI; data handling; prompting beyond templates.
LLM systems: retrieval, tool use, agents, function calling, memory, planning.
Evaluation: automated and human-in-the-loop evals; metrics; regression guards; observability.
Production: deployment, scaling, cost/perf tradeoffs, caching, safety and governance.

7 For Individuals

Efficient and systematic. Instead of scattered tutorials, you get a coherent pathway that compounds: concepts, exercises, projects, then capstones. You build a real portfolio and learn to ship maintainable AI features end-to-end. The result is practical fluency and employer-ready skills.

8 For Organizations

Build in-house capability—skip the talent wars. Rather than overpaying for scarce external hires, use a structured program to upskill engineers, data teams, and PMs together. The curriculum is designed for cohorts, enabling shared tooling, code standards, evaluation criteria, and deployment practices across teams.

9 Books Overview

The books are delivered as living HTML/PDF resources for both tracks. Each book comes with an accompanying GitHub repository for easy cloning and *one-click* opening of Jupyter Notebooks in Google Colab. Below are concise summaries mirroring the site.

9.1 Core Track

Python & Mathematics for ML & DL (∼300 pp.) Integrated, hands-on introduction teaching Python and essential math together (vectorization, linear algebra, calculus, optimization) applied immediately to ML/DL tasks.

Deep Learning Basics with PyTorch (∼250 pp.) Practical, end-to-end supervised learning: datasets, models, training loops, regularization, schedulers, and evaluation.

Building a Large Language Model from Scratch (∼200 pp.) Implement transformers step by step: tokenization, attention, architecture, training, sampling; reason about trade-offs and customization.

AI Agents & Automation (in progress) From tool use to orchestration: plan, reason, and interact with APIs and external systems while controlling cost, latency, and reliability.

9.2 Engineering Track

Python Primer for Data Science and Deep Learning (∼100 pp.) Environments, packaging, testing, project structure—remove tooling friction and build solid habits from day one.

Mathematics for Machine & Deep Learning (∼170 pp.) Rigorous, proof-aware treatment of core theory: linear algebra, multivariate calculus, convex/stochastic optimization, probability, information, kernels, numerics.

Mathematics Supplementary Resources (∼50 pp.) A concise Practitioner's Summary and Visual Guide to the main notions from the mathematics volume.

Software, ML & AI Engineering (∼100 pp.) Turn experiments into production systems: APIs, containers, CI/CD, data/feature pipelines and drift, RAG, agents, scaling/latency, observability and governance.

10 Why The AI Engineer

Approach and outcomes

- **Portfolio-driven** learning with interview-ready, GitHub-friendly capstones.
- **Do-first pedagogy**—every chapter leads to exercises, test projects, and capstones.
- **System orientation** beyond prompts: tools, retrieval, agents, evaluation, and operations.
- **Production focus** on reliability, latency, cost control, safety, and governance.
- **Cohort cadence** with networking, office hours, and accountability.

Resources and delivery

- **Colab-ready** notebooks with pinned dependencies and one-click launchers.
- **Tailored books** across both tracks (1,250+ pages combined) from fundamentals to production, with companion repositories.
- **Resource launchpad** plus schedule and practice plan provided to participants.
- **Engineering practices** baked in: testing, packaging, CI/CD, containerization, and observability.

11 By the Numbers

100%	Self-contained learning path
8	Core modules
4–8	Weeks duration (fast/medium track)
Portfolio	GitHub-ready deliverables
300+	Python scripts across books
100+	Jupyter notebooks
30,000+	Lines of code
1,250+	Documentation pages

Dr. Yves J. Hilpisch | taeieng@tpq.io | theaiengineer.dev
linktr.ee/dyjh | +49 15679 712305

October 6, 2025