The AI Engineer Program

Kickoff

Dr. Yves J. Hilpisch

The Python Quants GmbH

1. AI Frenzy
2. The Challenge
3. The AI Engineer
4. `yoctoGPT` Special Module
5. Who Benefits
6. Resources and Delivery
7. Why The AI Engineer
8. Enroll Now
9. Contact Information

Section 1

Investment, markets, and infrastructure signal unprecedented AI momentum.

AI Frenzy – Big Tech Capex

Capital expenditures for AI are surging across leading tech firms.

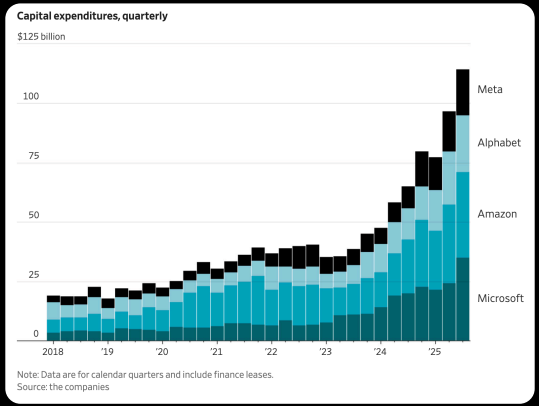

Source: WSJ – Big Tech Is Spending More Than Ever on AI

AI Frenzy – Compute Expansion

Google’s compute buildout underscores accelerating AI spending.

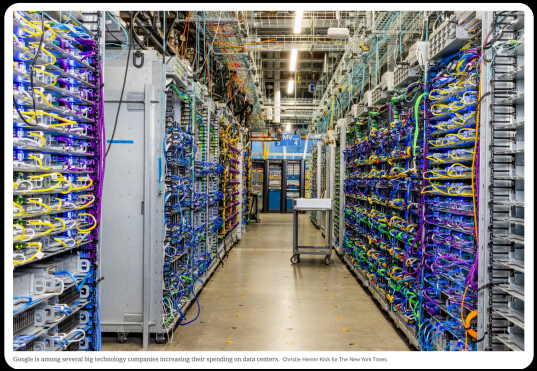

Google is among several big technology companies increasing their spending on data centers. Christie Hemm Klok for The New York Times

Source: NYTimes – AI Spending Is Accelerating

AI Frenzy – Market Expectations

Equity markets price in rapid AI-driven growth for platform leaders.

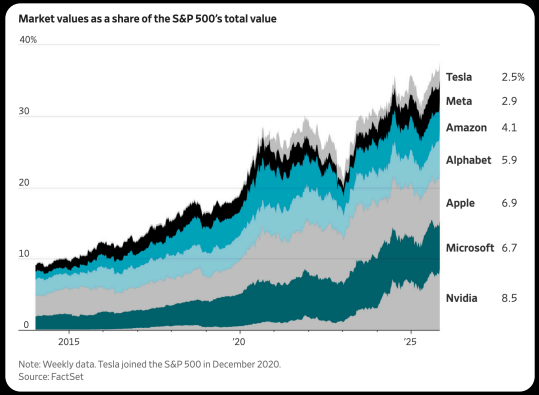

Source: WSJ – Big Tech Is Spending More Than Ever on AI

AI Frenzy – Infrastructure Boom

Nvidia’s valuation reflects the scale of compute fueling AI adoption.

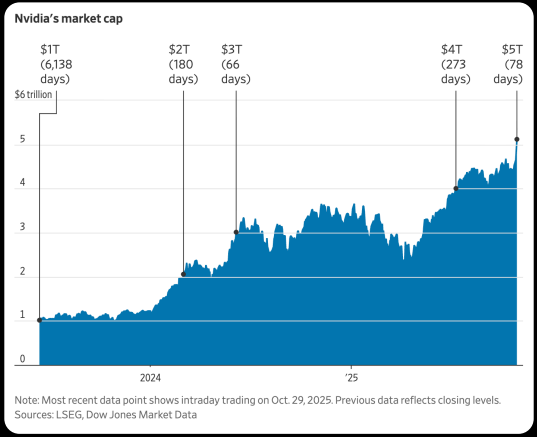

Source: WSJ – Nvidia First 5 Trillion Company by Market Cap

AI Frenzy – Efficiency Cuts Both Ways

Not everyone wins: companies push efficiency and junior job seekers face headwinds.

> *"OpenAI has more than 100 ex-investment bankers helping train its artificial intelligence on how to build financial models as it looks to replace the hours of grunt work performed by junior bankers across the industry."*
>
> *"The group, which includes former employees of JPMorgan Chase & Co., Morgan Stanley, and Goldman Sachs Group Inc., is part of a secretive project inside the startup that's code named Mercury, according to documents seen by Bloomberg."*
>
> *"Participants are paid $150 per hour to write prompts and build financial models for a range of transaction types, including restructurings and initial public offerings, according to a person familiar with the effort. The company has also granted the contractors early access to the AI it's creating that aims to replace entry-level tasks at investment banks."*

Source: Bloomberg – OpenAI Looks to Replace the Drudgery of Junior Bankers' Workload

AI Frenzy – Labor Market Impact

Surveys indicate AI-driven efficiency pressures affect entry-level hiring.

- **Headline finding:** 8-in-10 companies report plans to lay off recent college graduates due to AI adoption.
- **Drivers:** automation of entry-level tasks, efficiency mandates, and hiring freezes where roles are made redundant.
- **Impact:** fewer junior seats and a higher bar for applied AI fluency with portfolio evidence.
- **Implication:** candidates who can design, evaluate, and operate AI features are prioritized.

Source: Intelligent.com – 8-in-10 Companies Plan to Layoff Recent College Grads Due to AI

AI Frenzy – Seniority Gap

Early evidence shows GenAI complements senior roles and pressures junior ones.

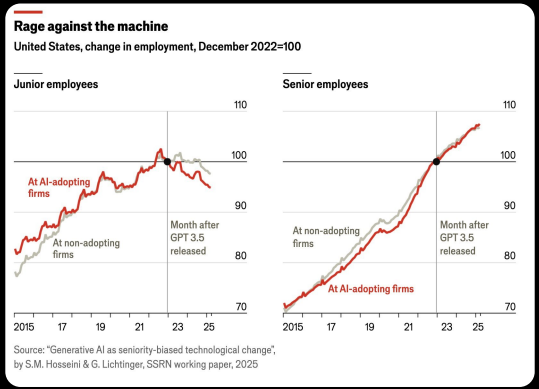

Source: SSRN – GenAI and Labor Market Seniority Effects

AI Frenzy – Talent Is A Premium Asset

In contrast, top AI engineering talent commands extraordinary compensation.

> *"So Mr. Zuckerberg personally met with Mr. Deitke. Then Meta returned with a revised offer of around $250 million over four years, with potentially up to $100 million of that to be paid in the first year, the people said. The compensation jump was so startling that Mr. Deitke asked his peers what to do. After many discussions, some of them urged him to take the deal – which he did."*
>
> *"Silicon Valley's A.I. talent wars have become so frenzied – and so outlandish – that they increasingly resemble the stratospheric market for N.B.A. stars."*
>
> *"Young A.I. researchers are being recruited as if they were Steph Curry or LeBron James, with nine-figure compensation packages structured to be paid out over several years. To navigate the froth, many of the 20-somethings have turned to unofficial agents and entourages to strategize. And they are playing hardball with the companies to get top dollar, much as basketball players shop for the best deals from teams."*

Source: NYTimes – AI Researchers Now Paid Like N.B.A. Stars

AI Frenzy – AI Reshapes the Workforce

AI deployment is capital-intensive, pressures parts of white-collar labor, and can simultaneously reward skilled workers who build the physical infrastructure behind the boom.

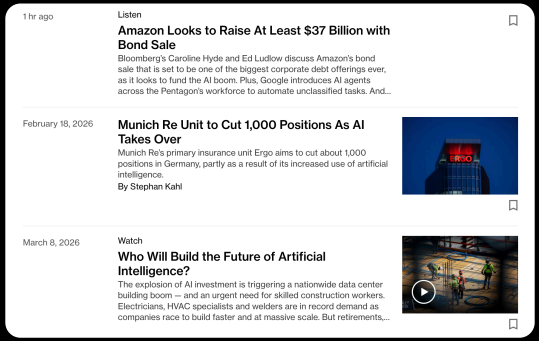

Source: Bloomberg, search for "ai workforce" on Mar 10, 2026

Section 2

AI reshapes work and creates pressure to build and ship.

The Challenge – AI Changes Work and Science

Modern AI reshapes work and increases demand for applied builders.

- **Automation of entry-level tasks:** chatbots, code copilots, and data copilots compress junior roles.
- **Demand shift:** teams need people who can *design*, *build*, *evaluate*, *and operate* AI systems.
- **Hiring reality:** organizations compete for applied AI talent; compensation reflects scarcity.
- **Implication:** a coherent, production-focused path beats scattered tutorials.

The Challenge – AI Use Becomes Mandatory

Leading tech firms increasingly move from encouraging AI use to enforcing it as a work expectation.

TECHNOLOGY • ARTIFICIAL INTELLIGENCE Follow

Tech Firms Aren't Just Encouraging Their Workers to Use AI. They're Enforcing It.

From startups to giants including Meta and Google, companies are factoring AI use into performance reviews and trying to track productivity gains

By Katherine Bindley Follow and Katherine Blunt Follow

Feb. 24, 2026 8:00 pm ET

Aa 102 Gift unlocked article Listen 〈7 min〉

Source: WSJ – Tech Firms Aren't Just Encouraging Their Workers to Use AI. They're Enforcing It.

The Challenge – AI Engineering Combines Disciplines

AI engineering blends established disciplines with emerging, LLM-era skills.

Established foundations

- Software Engineering (testing, packaging, APIs)
- Machine Learning Engineering (training, metrics, data)
- Data Engineering (pipelines, quality, governance)
- DevOps/SRE (CI/CD, containers, reliability)
- Security & Compliance (secrets, policies, reviews)
- Product & UX (requirements, feedback loops)

Brand-new or accelerated

- Prompting & System Prompt Design
- Retrieval & RAG (indexing, chunking, routing)
- LLMOps & Evaluation (ELO, rubrics, regression)
- Agents & Tool Use (planning, guardrails)
- Observability (traces, cost/latency, safety)
- Optimization (latency, cost, caching, batching)

Section 3

A production-first path from fundamentals to deployment.

The AI Engineer – From Zero to Production

Integrated tracks and outputs move participants from basics to deployment.

What you get

- **Dual-track pathway:** Core (ML/LLM/Agents) and Engineering (software/ML engineering).
- **Living books + repos:** 1,100+ pages across companion titles, Colab-ready notebooks, companion code, continuous updates.
- **Portfolio-driven**: exercises → test projects → capstones (interview-ready).
- **Production focus**: testing, packaging, CI/CD, containers, evaluation, observability, and safety.

Outcomes

- Ship end-to-end AI features and explain design choices.
- Build a public portfolio and speak to reliability, latency, and cost.
- Learn fast with coherent materials instead of piecemeal tutorials.

The AI Engineer – By the Numbers

The AI Engineer is backed by substantial, production oriented learning assets.

- **1,100+ pages**: living PDF books and summaries
- **8 modules**: four Core and four Engineering modules
- **7 companion books**: plus supplementary math resources
- **3 pacing options**: fast, medium, and self-paced
- **Portfolio**: interview-ready capstones

Delivery

Colab-ready notebooks, pinned environments, live orientation sessions, Discord support, and a resource launchpad.

The AI Engineer – Track Options

Three pacing options let delegates stay with one cohort while choosing the operating rhythm that fits their schedule.

Fast Track

4 weeks intensive

- going all in
- full-time commitment
- rapid skill acquisition

Medium Track

8 weeks balanced

- part-time compatible
- 2 weeks per module
- more practice time

Self-Paced

6 months self-paced

- flexible schedule
- full content access
- optional grading

The AI Engineer – Two Complementary Tracks

Two coordinated tracks cover core AI capabilities and engineering foundations.

Core Track

- ***Python & Math for ML/DL*** (vectorization, linear algebra, calculus, optimization)
- ***Deep Learning with PyTorch*** (datasets, training loops, schedulers, optimization)
- ***LLM from Scratch*** (tokenizers, attention mechanisms, GPT-style models)
- ***AI Agents & Automation*** (tool use, planning, guardrails, cost control)

Engineering Track

- Foundation Primers (Python best practices, math foundations, setup)
- Software Engineering (testing, packaging, CLI tools, CI/CD, containers)
- ML Engineering (data management, metrics, evaluation, reproducibility)
- AI Engineering (serving, latency, observability, safety, privacy)

The AI Engineer – Companion Books (7 Titles)

Concise, living books support both tracks with examples and exercises.

Core Track

- Python & Math for ML & DL (~188 pages)
- DL Basics with PyTorch (~251 pages)
- Building an LLM from Scratch (~183 pages)
- AI Agents & Automation (~106 pages)

Engineering Track

- Python Primer for DS & DL (~100 pages)
- Mathematics for ML & DL (~162 pages)
- Mathematics Supplementary Resources (~48 pages)
- Software, ML & AI Engineering (~93 pages)

Book Samples: theaiengineer.dev/#books

The AI Engineer – Schematic Overview

A concise map of how tracks, content, practice, and portfolio elements align.

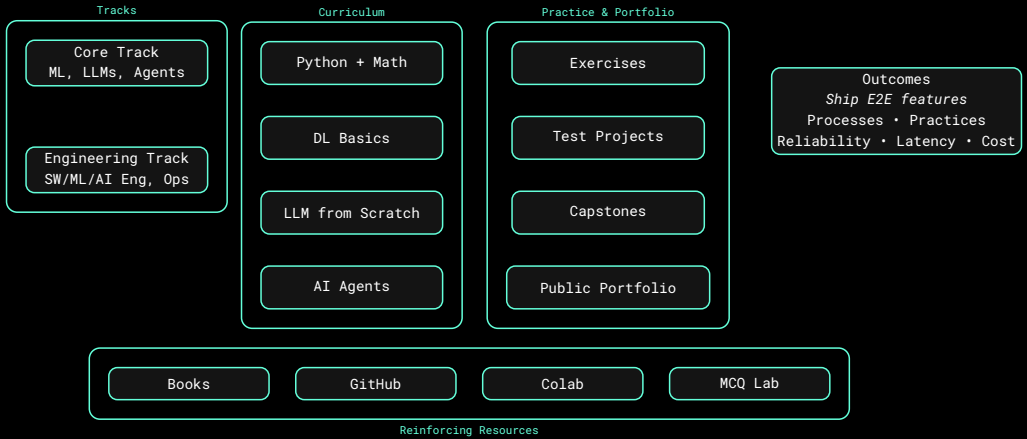

The AI Engineer – Portfolio Examples

Four main capstones anchor the delegate portfolio and mark the key Core track milestones.

- **GD/SGD Optimization milestone:** build and analyze gradient-based optimization workflows from calculus to implementation.
- **NN Backpropagation milestone:** implement a small neural network and make the chain rule and backpropagation fully explicit.
- **Attention/Transformer milestone:** build a tiny transformer-style model and connect attention math to executable PyTorch code.
- **MCP Agents milestone:** design an MCP-based agent workflow with tools, orchestration, observability, and deployment discipline.
- **Completion signal:** a TPQ certificate may be issued for successful completion based on portfolio review against typical industry practice expectations.

Section 4

A new special module extends the program with realistic GPT building and training.

yoctoGPT Special Module – What yoctoGPT Adds

yoctoGPT adds a compact but realistic build-from-scratch GPT/LLM module to the program.

- It focuses on implementing a real GPT-style language model from first principles instead of only using high-level APIs.
- It includes full supporting documentation and structured notebooks for small character corpora, larger corpora, token-based training, and finance-oriented token corpora.
- It is designed around practical training workflows for Google Colab TPU and GPU instances.
- It makes model building, preprocessing, tokenization, training, sampling, and optimization visible end to end.

yoctoGPT Special Module – yoctoGPT Learning and Execution Flow

The yoctoGPT documentation teaches the system as a staged flow from mental model to reliable Colab execution.

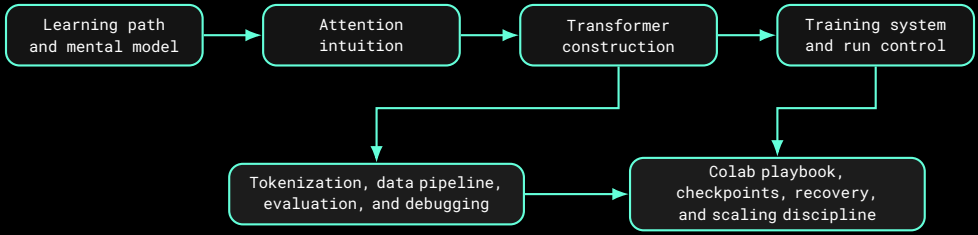

Adapted from the yoctoGPT documentation: learning path, training system, and Colab playbook chapters.

yoctoGPT Special Module – Why yoctoGPT Fits The AI Engineer

yoctoGPT is a natural extension because it connects the program's Python, math, deep learning, and LLM material in one executable system.

- It turns transformer concepts into a full training pipeline that participants can inspect, modify, and run.
- It strengthens the bridge from *LLM from Scratch* to hands-on engineering on realistic hardware budgets.
- It gives advanced learners a new portfolio artifact that demonstrates deeper model understanding than prompt-only work.
- As a brand-new special module, it expands TAE with an explicit path from fundamentals to custom GPT training in Colab.

Section 5

Students, academics, and professionals turn learning into portfolios.

Who Benefits – Students and Academics

Academic learners gain a bridge from theory to deployable, evaluated projects.

- **Students:** move from coursework to deployable projects; portfolio gives hiring signal.
- **Academics:** translate research into robust demos; evaluation harnesses reduce friction.
- **Why now:** AI fluency is the new baseline across disciplines; credibly shipping matters.

Who Benefits – Professionals and Corporations

Working professionals and teams adopt shared practices to ship reliable AI.

- **Professionals:** upskill to applied AI with guardrails; learn patterns you can explain on the job.
- **Corporations:** shared standards for teams – testing, packaging, CI/CD, evaluation, governance – across software, ML, and **AI engineering**.
- **Why now:** production AI is moving from pilots to owned capability; recruiting is scarce.

Section 6

Living materials, reproducible repos, and deliberate practice cadence.

Resources and Delivery – Just a Few Clicks

Materials and delivery emphasize speed, reproducibility, and deliberate practice.

- **Living books:** PDF, 1,100+ pages, updated and aligned; rich visuals and examples.
- **Repos + Colab:** one-click notebooks; pinned environments; reproducible exercises.
- **Track options:** fast, medium, and self-paced learning paths.
- **Support:** Discord, office hours, launchpad, and practice schedule.

Resources and Delivery – Launchpad

A single hub links books, repos, Colab notebooks, and practice.

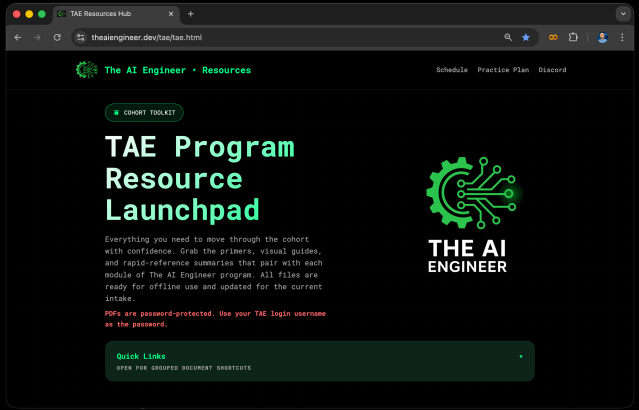

Source: theaiengineer.dev/tae/tae.html

Resources and Delivery – Schedule

A cohort schedule outlines weekly focus, deliverables, and time planning.

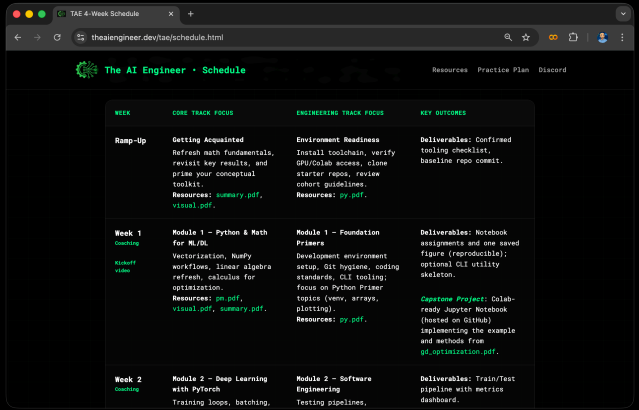

Source: theaiengineer.dev/tae/schedule.html

Resources and Delivery — Python & Mathematics Book

The Python & Mathematics for Data Science and Machine Learning book supports the Core track foundations.

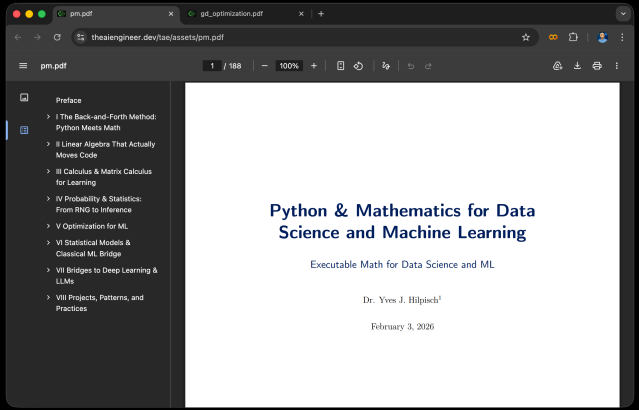

Resources and Delivery – GD Optimization

Gradient-based optimization notes build the basis for the Week 1 capstone project.

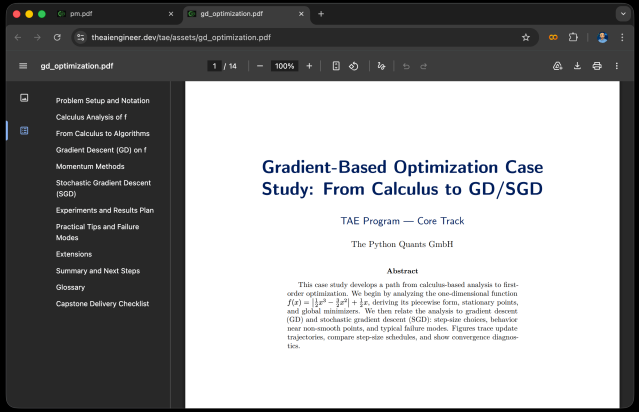

Section 7

Portfolio-driven approach ties methods to measurable outcomes.

Why The AI Engineer – Approach and Resources

A portfolio driven, production mindset links methods to measurable outcomes.

Approach & outcomes

- Portfolio-driven, interview-ready capstones.
- Do-first pedagogy: exercises → projects → capstones.
- Skill acquisition: detailed steps stay in the self-paced materials.
- System orientation: tools, retrieval, agents, evals, ops.
- Production focus: reliability, latency, cost, safety.

Resources & delivery

- Tailored books and companion repos.
- Colab-ready notebooks; one-click launchers.
- Live + recorded sessions, Discord support, and office hours.
- Resource launchpad, schedule, and practice plan.
- A TPQ certificate may be issued after successful portfolio review.

Section 8

Join the cohort to build and showcase real AI work.

Enroll Now – Special Offers and Enrollment Options

Enrollment stays simple, and the website carries the current cohort details and checkout information.

- Introductory and special offers may be available; check the website for current details.
- Join the program and start right away with guided study suggestions.
- Learn more and enroll: theaiengineer.dev
- Bundle option with Python for Finance (CPF) and The Crypto Engineer (TCE).
- Learn about the bundle: python-for-finance.com

Dr. Yves J. Hilpisch

The AI Engineer

TAE: theaiengineer.dev
Launchpad: theaiengineer.dev
Email: team@tpq.io
Linktree: linktr.ee/dyjh