tripleten

AI & Machine Learning

For tech professionals with production experience

20-week accelerator

Online

10–15 hours a week

How this program works

Learn to build senior-level AI systems, and make your next career step

5 projects · 1 capstone you design

10 h a week on projects · about 16 with all the theory

20 weeks

A year of career support

1

Five production projects, five live defenses

Five projects from real briefs, each with a pre-defense and a live defense in front of a working engineer. Theory is optional: skip what you know. 10 hours a week on projects.

2

Work with a senior engineer from day one

Instructors from AMD, Microsoft, Amazon and Capital One. A 30-minute 1:1 every week, written feedback on every project, your project space always on.

3

Career prep for every round, and for the search itself

Up to 10 live mocks: screening, technical, recruiter and behavioral, with a plan after each. A coach rewrites your resume and LinkedIn, the platform matches roles and tailors your resume to each posting. Support for up to 12 months.

Program summary

Five production projects, defended live

Five projects, each a real task companies solve in production: a product requirement, a codebase or a dataset, and acceptance criteria known upfront. Not a sandbox. Projects 1 and 2 are one roadmap: the model you train in Project 1 is the one you ship in Project 2. From Project 3 on, each brief is a new company; the capstone is an open brief on an existing product.

~200 h of project work · ~125 h of theory | **20 weeks total**

Project 1	Train and track an ML model on real-world data	~95 h
Project 2	Deploy the model as a production service	~57.5 h
Project 3	Fine-tune and evaluate a clinical LLM	~60 h
Project 4	Build a retrieval agent and harden it	~60 h
Project 5	Add AI features to an existing product	~52.5 h

Every project ends with a live defense

At about 70% you book a 30-minute pre-defense with your instructor, who has read your repository in advance and flags structural problems while there is still time to fix them. When the project is done, you defend it in a 60-minute live defense with a working engineer. If it doesn’t pass, you get written feedback and book again.

How long will it take you?

Project hours are fixed and set the timeline. Theory is optional and flexes with your experience, so two people can finish the same program on different schedules.

5 h a week	**10 h a week**	15 h a week	20 h a week
40 weeks on projects only	20 weeks on projects only	14 weeks on projects only	10 weeks on projects only
65 weeks with all the theory	33 weeks with all the theory	22 weeks with all the theory	17 weeks with all the theory

Based on 200 hours of project work and 325 hours with all the theory. Plan for 10 hours a week on projects. From there, how long it takes depends on the hours you invest each week, your starting level and how engaged you stay. The program is flexible and fits around your current schedule.

Project 1

Train and track an ML model on real-world data

You join Wayfinder Labs and turn a product requirements doc into a working model: pull data from the MBTA's live API into a database, then train, evaluate and track a model that predicts whether a bus will be late. You finish with a model that solves the stated business problem, versioned and registered the way a real company would ship it.

Project work	~50 h
Theory (optional)	~45 h
Total	~95 h

Theory topics

- Project brief: from PRD to ML spec
- Claude Code for AI engineering
- The web as a data source; APIs to database: the ingestion pipeline
- Regression and classification fundamentals; training in scikit-learn
- Preparing features; encoding and scaling
- Accuracy and the confusion matrix; precision, recall and class imbalance; thresholds and ROC/PR curves; regression metrics
- How neural networks learn; PyTorch fundamentals; training deep networks well
- Experiment tracking with MLflow

Skills you'll add

Classification **PyTorch** **MLflow & DVC** **Data ingestion pipelines** **Claude Code**

Learning objectives

Spec & baselines

Translate a PRD into an ML spec with framing, success metrics and baselines.

Ingestion pipeline

Build an ingestion pipeline from a live API through validation into PostgreSQL.

Train & evaluate models

Train and evaluate classification models and choose the right metrics.

Deep learning done right

Build and train PyTorch models with sound training mechanics, tracked in MLflow.

Project 2

Deploy the model as a production service

You take the Pulse model out of the notebook and run it the way production runs: a FastAPI service in Docker, two replicas behind a proxy with health checks and restart policies, drills that break it on purpose, the whole stack described as code and checked on every push. Everything runs on your laptop and a GitHub repository. The cloud is one optional page.

Project work	~30 h
Theory (optional)	~27.5 h
Total	~57.5 h

Theory topics

- FastAPI for ML services; model serialization and management; API best practices
- Docker fundamentals; building Docker images; Docker for ML applications
- OpenTofu for local infrastructure
- MLOps foundations and testing; observability and CI

Skills you'll add

FastAPI · Docker & Compose · OpenTofu · GitHub Actions · Failure drills & runbooks

Learning objectives



API serving

Serve a PyTorch model with FastAPI and Pydantic, tested against the real model.



Run it like production

Run a multi-container stack with replicas, a reverse proxy, health checks and restart policies.



Infrastructure as code

Describe the stack in OpenTofu and prove plan, apply, destroy and rollback.



Gate changes

Run lint, tests and an infrastructure plan on every push.

Project 3

Fine-tune and evaluate a clinical LLM

You teach an open LLM to write clinical visit summaries the way hospital doctors write them. You curate and de-identify the training data, run parameter-efficient fine-tuning, and evaluate the generated text against a clinician rubric rather than a loss curve. A generic chatbot can’t be trusted with this, which is the point.

Project work	~40 h
Theory (optional)	~20 h
Total	~60 h

Theory topics

- Transformers and BERT; LLM fundamentals; generation parameters
- PEFT and LoRA; QLoRA
- Hugging Face fine-tuning workflow
- Clinical data curation and de-identification
- GPU training
- LLM evaluation

Skills you’ll add

LLMs **PEFT, LoRA & QLoRA** **Data curation & PHI discipline** **Hugging Face** **LLM evaluation**

Project 3

Learning objectives

Transformers & attention

Explain transformer architecture, attention and the BERT-to-LLM lineage.

Fine-tuning (LoRA/QLoRA)

Apply parameter-efficient fine-tuning: adapters, LoRA and QLoRA.

Data curation

Curate instruction datasets for a regulated domain with de-identification.

Honest evaluation

Evaluate generated clinical text with automatic metrics, factuality checks and a rubric.

Build a retrieval agent and harden it

You build Resolve, an AI system that does the first pass on a software company's support tickets: it screens each ticket for attacks, classifies it, sends it to a specialist agent that researches the answer in the company's documentation and acts in the helpdesk, has a second agent review the draft, and then resolves the ticket or hands it to a human with the research attached. You prove it is safe by attacking it yourself, and prove it is good by measuring it on tickets it never saw. Everything runs on your laptop.

Project work	~40 h
Theory (optional)	~20 h
Total	~60 h

Theory topics

- Prompt engineering; calling a hosted model from Python
- Retrieval and RAG with OpenSearch; RAG warm-up build
- Agent fundamentals; LangChain and LangGraph; specialists and a reviewer
- MCP: tools for agents; running the agent as a service
- LangSmith observability
- LLM-as-judge and hallucination metrics
- OWASP threat modeling; defenses and red-teaming

Skills you'll add

RAG & OpenSearch · LangGraph · MCP · Guardrails & prompt injection · OWASP for LLMs · LLM-as-a-judge · LangSmith

Project 4

Learning objectives

Retrieval with numbers

Build a retrieval pipeline and measure recall on dev and held-out sets.

Multi-agent workflow

Design a workflow with specialists, a reviewer, guards and decisions in code, traced in LangSmith.

Tools through MCP

Write and consume an MCP server with least privilege and idempotent writes.

Attack your own system

Threat-model with the OWASP LLM Top 10, build an injection suite and red-team it.

Project 5

Add AI features to an existing product

The capstone. You join the team behind Atlas, a healthcare platform that already exists. You clone the provided app, run it locally on Kubernetes with MongoDB, explore the sample patient data, and get the brief AI engineers get in industry: “Here’s our product and our data. What intelligent features can you add?” You train two predictive models, in scikit-learn and PyTorch, on AWS GPUs through the Codio cloud lab, export them with ONNX, and add one generative or agentic feature. The platform is the stage; your AI work is the show.

Project work	~40 h
Theory (optional)	~12.5 h
Total	~52.5 h

Theory topics

- Scoping AI features against an existing product: feasibility, value, deterministic versus generative
- Integrating AI into SaaS products
- Handling patient-style data responsibly in a prototype
- Demo preparation and product rationale
- Running the product stack locally: Kubernetes and MongoDB
- GPU training on AWS through Codio; ONNX export

Skills you’ll add

Feature scoping · SaaS integration · Local Kubernetes · MongoDB · GPU training on AWS (Codio) · ONNX · Product demo & defense

Project 5

What the final system does

Atlas with your features inside. Two predictive models trained on the product's own data surface risk scores in the app, and one generative or agentic feature turns a patient record into insights. You demo it live and defend the product decisions with your instructor.

Learning objectives

Feature scoping

Scope AI feature ideas against an existing product and make the deterministic-versus-generative call.

SaaS integration

Integrate new services with a provided containerized application.

Predictive features

Build and train predictive models on provided data, on AWS GPUs through Codio, and expose them as product features.

Demo and defend

Demo working prototypes and defend the product decisions behind them.

Who will teach you

Working engineers, not career lecturers

Your instructors build production ML systems for a living. They meet you one-on-one every week, run your project defenses and your technical mock interviews. More experts join against the same bar.



Instructor

Gagandeep Singh

Capital One · Walmart · Chime

GitHub LinkedIn

AI/ML engineering lead with 10+ years shipping production ML: recommendation systems, personalization and LLM platforms; leads teams that put agentic systems into production.



Instructor

Mohamed Cherif

AMD · Microsoft · Amazon · Skype

GitHub LinkedIn

Principal-level engineer with 20+ years building GPU-accelerated media pipelines, SDKs and computer vision systems.



Program author

Amish Suchak

Xsolis · MIT Professional Education · AI4ALL

GitHub LinkedIn

Data scientist and ML engineer with 10+ years building and deploying models to production. Co-founded Percruit, the career platform you'll use in the program.

Coaches who have been on the hiring side

Your career coach joins after your first project and stays with you for up to 12 months: target role, resume and LinkedIn, recruiter and behavioral mock interviews, and a plan after every real interview. Technical mocks are run by the program's instructors.



Career coach

Monica Bencze

Google

Flatiron School

LinkedIn

Recruited for Google. Knows what a top tech company screens for.



Career coach

Francesca Ogilvie

Deloitte UK

DLA Piper

LinkedIn

Ran 40+ assessment centers at Deloitte UK. Knows how hiring decisions get made from the inside.



Career coach

Peter Spellman

Berklee College of Music

Boston University

LinkedIn

30+ years in career development. Led Berklee's career center for 24 years.



Career coach

Daniela Ruiz

Microverse

LinkedIn

Coached 800+ software development bootcamp students toward jobs.

Your career track starts after Project 1 and runs for up to 12 months

We reposition the career capital you already have and get you ready for every interview round. Your career coach joins after your first project and stays with you for up to 12 months.

1

Positioning & target role

Your coach joins after your first project: career platform access, target role and a transition plan.

2

Portfolio & materials

Every defended project becomes a case study. Resume and LinkedIn rewritten for the target role, and tailored to each posting on the career platform.

3

Mock interviews & debriefs

Up to 10 live mocks for every round with instructors and your coach, unlimited AI practice, a plan after each mock, and debriefs of real interviews.

4

Job search & a year of support

Partner job openings and role matching on the career platform. About 4 months of active career work, then up to 8 months while you go through your transition.

A cohort of engineers at your level

You study next to working engineers from different companies and backgrounds. The median student has 5 years in the profession, and some bring 20 years or more.

How the community works

- **Community hub from day one**
  Announcements, events and a member directory with LinkedIn and GitHub. Your access stays after the program.
- **Your project space**
  You're always next to people working on the same problem. Instructors answer there between sessions.
- **Moderated networking events**
  Regular events and guest speakers from the industry.
- **Random coffee**
  We pair you with another student for a one-on-one chat.
- **TripleTen alumni**
  After your capstone you join a network of 7,500+ graduates.

Who you'll study with

Our students have worked at

Chase · Wells Fargo · PNC · KeyBank · Meta · Amazon · Apple · Intel · HP Enterprise · NASA · Northrop Grumman

Backgrounds: software engineering, DevOps and cloud, data engineering and BI, ML and AI engineering, security, QA.

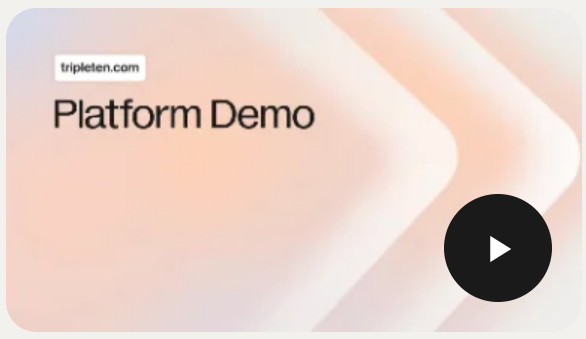
Names are kept private to protect our students' confidentiality. All live activity runs in the evenings and on weekends.

A typical 20 weeks

At 10 hours a week on projects. The career track runs alongside your projects: its milestones are the highlighted rows. Theory is optional and flexes with your experience; week numbers are guides, not deadlines.

Before start	Start	A call with your onboarding manager to build your personal study plan
Day 1	Start	Orientation and your first project brief. You join the community hub and meet your cohort
Weeks 1–5	Project 1	Train and track an ML model on real-world data
~Week 5	Career	Your career coach joins: career platform access, target role and transition plan
Weeks 6–8	Project 2	Deploy the model as a production service
Weeks 9–12	Project 3	Fine-tune and evaluate a clinical LLM
~Week 12	Career	Resume and LinkedIn rewritten for the target role; defended projects become portfolio cases
Weeks 13–16	Project 4	Build a retrieval agent and harden it
Weeks 13–20	Career	Mock interviews for every round: up to 10 live sessions with instructors and your coach, unlimited AI practice
Weeks 17–20	Project 5	Add AI features to an existing product
Graduation	Career	Capstone demoed live and defended with your instructor. Career support continues for up to 12 months from the start of the track

Take a look inside the program before you commit



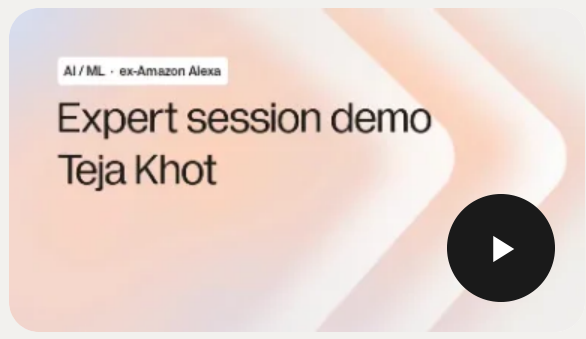
Platform demo

Three minutes inside the platform: your project map, the skill-gap check, theory you can skip, booking a 1:1, pre-defense and defense, events and the community hub.



Program walkthrough

The program's expert walks through the five projects, how defenses work and what you finish with.



A live session

A five-minute fragment of a group office hours session with an instructor: how they work with students.

Comparing your options

Most options give you either learning or career support. This program gives you both, at a senior level.

	This program	Self-study	Online courses	Master's and university programs	Interview prep services
Duration	20 weeks	Flexible	Hours to weeks	1–3 years	3–8 months
Production projects, defended live	✓	✗	✗	✗	✗
Weekly 1:1 with a senior engineer	✓	✗	✗	✗	✗
Mock interviews	✓	✗	✗	✗	✓
Resume and LinkedIn repositioning	✓	✗	✗	✓	✓
Fits around a full-time job	✓	✓	✓	✓	✓

Answered straight

Is this live classes or self-paced?

Neither, and that’s the point. You study on your own schedule, with no mandatory class times. But it isn’t a self-paced course library: every week there’s a 30-minute 1:1 with an instructor and live group sessions, and every project ends with a live defense in front of a working engineer.

Do I have to quit my job?

No. The program is designed for people working 40 hours a week or more. There are no fixed class times, and all live sessions run in the evenings and on weekends. Plan for 10 hours a week on projects, and up to 16 if you take all the theory. At fewer hours a week, plan for longer.

Am I qualified? My Python is not strong.

This program needs Python as a working language from the first project: the whole stack runs on it, from scikit-learn and PyTorch to FastAPI and LangGraph. If you write code in another language and your Python is rusty, the on-ramp has refreshers for git, SQL and pandas, but not a Python course. If you don’t write Python at all, AI Systems Engineering is the better fit.

I’m already advanced in my career. Will this teach me anything new?

That depends entirely on where your gaps are. It is why the call starts with the syllabus and how it compares to your current skills and experience. If you partially know the material, you’ll be able to skip the theory lessons and focus on building.

Who actually teaches, a person or AI?

People. Our instructors run the one-on-one sessions and the live project reviews, and someone reviews your pull request and decides whether you are ready to defend it. AI tools are part of what you learn to build, but you’ll learn from experienced professionals.

How much instructor time do I actually get?

A weekly 30-minute 1:1, written feedback on every project, a pre-defense and a live defense per project, and your project space where instructors answer between sessions.

Will I study alone, or is there a real cohort?

You move in project cohorts grouped by the project you're on, with a project space, live events, a member directory and coffee pairings. Working engineers around you, not a beginners room.

Is TripleTen accredited?

TripleTen is licensed by the Wyoming Department of Education and is registered under California law as an out-of-state institution providing private postsecondary education. You finish with a completion certificate and a defended portfolio of production work.

The program is brand new. How do I know it works?

These topics are new, but the program is built on the same foundation as our established programs: the same platform, project review process and career support. The curriculum is built and approved by engineers and hiring managers from AMD, Microsoft, Amazon, Capital One and Xsolis. To see if it's the right fit, review the instructors' profiles, watch the program walkthrough and read this syllabus module by module. And the first 14 days are fully refundable, so you can start, meet your instructor and see the platform before you commit.

Is there a job guarantee?

The hiring decision is always the employer's, so no program can guarantee you a job, and this one has no money-back guarantee. What we commit to is everything in our control: the skills the market is hiring for right now, production projects defended in front of working engineers that show the quality of your work, and career support throughout the program, ensuring that you are up to date with the market trends.

Can I try it before I commit?

Yes. The first 14 days are fully refundable. You can start the program, see the platform and meet your instructor, and still leave at no cost if it's not what you expected.

Will I have to start my career over?

No. This program builds on the experience you already have and prepares you for your next step, not a reset. The curriculum prepares you for senior-level work, and our career team leverages your existing experience as a key qualification. While we ensure hiring managers clearly see the value of your background, final placement and title decisions ultimately rest with the individual employer.

Who this program is for

A quick self-check before you commit. Honest in both directions.

You're a good fit

- You have 2+ years shipping production code or working with data: software engineering, data engineering, analytics, data science, ML, DevOps, QA or quantitative work. You write Python.
- You already use AI at work but haven’t built AI systems yet, and you feel that gap in interviews or promotion conversations.
- You have a strong tech base but under 2 years in production: a junior engineer, analyst, QC or IT specialist, or a CS degree, master’s or PhD.

This is probably not for you

- You’re starting from zero, with no coding experience yet. TripleTen’s core bootcamps are the right entry point; this program assumes a working base.
- You’re after a certificate rather than a portfolio of defended work.

Two or more years is a guide, not a cut-off. Someone hands-on in the last six months is in a stronger position than someone whose last production work was eight years ago.

What you need on day one, and what's included

Prerequisites

Skill	What we expect
Python	A working language from Project 1: functions and classes, comprehensions, try/except, imports and packages, virtual environments and pip, reading and writing files, reading a traceback
Command line	Run commands in a shell
Git	Clone, commit, branch, push
SQL	Basic: SELECT, WHERE, GROUP BY, simple JOINs. Project 1 loads data into Postgres and queries it from the first task
pandas	Basic: load a CSV, filter, group, merge
Statistics	Enough to read metrics: averages, distributions, why a train/test split exists. No calculus or linear algebra; the ML math is taught in the projects

Taught inside, not needed on entry: scikit-learn and PyTorch, Docker and Docker Compose, FastAPI, infrastructure as code and CI, fine-tuning, RAG and agents, LLM APIs, evaluation, Kubernetes and MongoDB at capstone level. The on-ramp before Project 1 has refreshers for git, SQL and pandas. No Python at all? AI Systems Engineering is the better fit.

Your computer

16 GB RAM · 4 cores · macOS on Apple Silicon, Windows or Linux · admin rights to install Docker Desktop · no GPU on the laptop.

Use a personal machine: work laptops often block installs.

Included in your tuition

- A Claude Pro subscription for the length of the program, on request
- AWS GPUs through the Codio cloud lab for Project 5, no cloud account on your side. Project 3 fine-tuning runs on your own machine or in free Colab
- Model API access for the agent project through the program's provider organization
- Access to your track's course material after the program, kept up to date

Ready to ship your AI product?

TripleTen is a career learning platform
for a world being reshaped by AI.

Book a call to talk to a career mentor

tripleten