

USE IT IN YOUR WORK

Working in AI

The roles that exist, what they pay, and how people get in without a famous degree.

First edition, September 2026

Addaly

Series editor: Dr Mustafa Mun

Draft, open for academic review

USE IT IN YOUR WORK

Working in AI

The roles that exist, what they pay, and how people get in without a famous degree.

Series editor: Dr Mustafa Mun

Addaly

First edition, September 2026 · Draft, open for academic review

Working in AI

© 2026 Addaly.

Licensed under Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0). You may copy, print, share and adapt it for any non-commercial purpose, with credit, under the same licence.
creativecommons.org/licenses/by-nc-sa/4.0

First edition, September 2026, build 62a26075.

Written by AI models to a written standard, with Dr Mustafa Mun as series editor. Exam keys checked blind by a second model: 3,665 of 3,666 agreed. Academic review invited.

The manuscript was drafted with the assistance of a large language model and was edited and published under his name. That is stated plainly here rather than left to be discovered, because a reader is entitled to know how a book they are trusting was made.

How to cite: *Mun, M. (2026). Working in AI. Addaly. addaly.com/learn/ai-careers.*

This book is the printed form of the course of the same name at addaly.com/learn/ai-careers. The website is the living version and is corrected first. Where the two differ, the website is right.

Free to read, free to download, free to print, and free to teach from. There is no paid edition and there never will be.

Every diagram in it is drawn from data by code, not generated as a picture, so the numbers on the page are the numbers in the argument. The answers to every check, test and examination question are at the back.

9 modules · 85 lessons · 33 figures · 9 case studies · 65,131 words · about 12 hours

Brief contents

	Contents	6
1	The terrain: who hires, for what, and where you fit	11
2	The roles, one at a time	59
3	The evidence: what gets tested, and how to prove you have it	107
4	Getting seen: the routes that put a person in front of your work	167
5	Using AI in your own search	215
6	The interview, from first screen to signed offer	259
7	The money, and the paperwork underneath it	311
8	Freelancing and contracting, as a business	355
9	Staying: the first ninety days and the ten years after	397
	Examination	443
	Answers	467

1

MODULE 1

The terrain: who hires, for what, and where you fit

Name the employer types that hire for AI work in your own market, decode a job advert into the role it really is, and commit to one lane with a plan sized to the hours you actually have.

1	The jobs, described as they feel on a Tuesday	12
2	The eight kinds of employer, and how many of each	15
3	Geography, time zones, and the office you have never heard of	20
4	The work that trains the models	24
5	Reading a job advert for what it actually is	27
6	Who sits around you, and how work arrives	31
7	Who actually needs a degree	35
8	Choosing a lane you can defend, and changing it later	38
9	Levels, titles, and why "senior" means nothing on its own	42
10	A year on six hours a week	46
	<i>Case study: Two offers, one title, and the question about the dashboard</i>	51
	<i>Module 1 test</i>	56

Lesson 1 · 7 min

The jobs, described as they feel on a Tuesday

THE ONE THING TO KEEP

Most AI work is moving messy data and checking outputs, not training models.

Titles lie, work does not

"AI Engineer" means at least six different jobs depending on who wrote the posting. Read the responsibilities section, ignore the title. Here are the roles that actually exist, described as they feel on an ordinary Tuesday rather than in a brochure.

Data engineer

You move data from where it is created to where someone can use it. Pipelines, schemas, backfills, an alert at 03:40 because last night's job died on a row containing a comma.

A Tuesday: a dashboard shows Monday's orders as zero. You trace it to a payments provider that renamed a field from `order_id` to `orderId` on Saturday without telling anyone. You fix the parser, replay two days of data, and write a note so the next person loses ten minutes instead of three hours.

There are more of these jobs than any other kind on this list, in every country. Least discussed, most stable.

Data analyst

SQL, spreadsheets, dashboards, and the question behind the question. Someone asks for "signups by region". What they are actually deciding is whether to keep paying for the Nairobi campaign. Most of the skill is finding that out before you write the query, because the query takes four minutes and the wrong question wastes a week.

This is the most common first job in the field for people who did not study computer science.

ML engineer

You put models into products and keep them alive: serving, latency, retraining, monitoring, rolling back when the new version is quietly worse. Training models is a slice of it. Training one from scratch is rare outside research labs.

A Tuesday: recommendations got 400ms slower on Friday and nobody knows why. You find a feature lookup making one database call per item instead of one per request.

AI product manager

You decide what gets built and what "good enough" means — which in this field means writing down how the system will be measured before it exists. Also: arguing with legal about which data may be used, cutting scope, and telling people no in a way they can accept.

Prompt and evaluation work

The newest and least settled corner. The durable half is not writing prompts, it is evaluation: assembling the test cases that decide whether a system ships, grading outputs, hunting failure modes, turning "this feels off" into a number someone can act on. Job titles here change every eighteen months. The work does not.

Research scientist


Reading papers, running experiments that mostly fail, writing up the few that do not. A small number of jobs, concentrated in a few dozen labs and universities worldwide, nearly all wanting a PhD. Keep it on your map. Do not plan your life around it.

The ratio nobody prints

Picture a 300-person fintech in São Paulo with a genuine data team: about four data engineers, six analysts, two ML engineers, one product person who covers AI alongside three other things, and zero researchers. Multiply that shape across thousands of companies and you have the actual job market.

Figure 1.1

One real data team, at a 300-person fintech



If you are choosing where to aim, aim where the jobs are and where you can produce evidence fastest. Analyst and data engineer are the two widest doors. ML engineering is far easier to reach from a data engineering job than from a course.

What they all share

Every role above is mostly one activity: something is wrong with the data and you have to find out why. Courses skip this because it does not demo well. Employers pay for it because it is where the time goes.

BEFORE YOU MOVE ON

A 300-person logistics company in Jakarta hires its first AI person to build an assistant that answers staff questions about shipments. Three months in, where has most of that person's time gone?

- A. Fine-tuning a language model on the company's shipment history.
- B. Getting shipment data out of four systems that disagree with each other, and building a set of test questions to check the answers against.
- C. Comparing model providers and negotiating a contract with one of them.
- D. Rewriting the prompt until the answers read well.

The answer and why: p. 469

Lesson 2 · 8 min

The eight kinds of employer, and how many of each

THE ONE THING TO KEEP

Nearly every AI-adjacent job sits inside an ordinary organisation that uses software, not inside a company that builds models.

The picture in most people's heads is the smallest box

Ask someone where AI jobs are and they name four or five companies that train frontier models. Those organisations together employ a few thousand research staff worldwide, nearly all with doctorates, nearly all in a handful of cities. That is a real employer category. It is also the one you are least likely to work for and the one that shapes almost none of the market you will actually apply into.

Here is the fuller list. Read it as a map of where to spend your applications.

1. Frontier labs and big-tech research

Tiny, famous, PhD-heavy, publication-driven. Salaries at the extreme top. If you are not already publishing, this is not a plan.

2. Big-tech product teams

Large and real — the recommendation, search, ads, fraud and infrastructure teams. They hire steadily but skew senior, and their junior intake mostly comes through internships and graduate programmes with fixed windows. If you are outside that pipeline, entry is hard and lateral entry at three years' experience is much easier than entry at zero.

3. IT services and outsourcing

Infosys, TCS, Wipro, HCL, Cognizant, Accenture, Capgemini, Globant, and hundreds of smaller firms. In India, the Philippines, parts of Latin America and Eastern Europe this is the single largest employer of junior technical people, by a wide margin.

The economics are worth understanding because they change your strategy. These firms bill clients for your time and often bill more for certified staff, which is why a cloud certification has genuine commercial value here and almost none elsewhere. You may sit "on the bench" between projects. Training is real but generic. Pay starts low and rises slowly.

The honest read: it is a legitimate and very wide door, and a large number of working data engineers walked through it. Take the job, get two years, then move.

4. Product startups

Twenty to three hundred people. They read applications, they hire on evidence, and there is no qualifications filter because there is no HR department to run one. This is where a portfolio does the most work per hour you put into it. Against that: less mentorship, more chaos, and a real chance the company is gone in eighteen months.

5. Enterprises with internal data teams

Banks, insurers, telecoms, retailers, airlines, hospitals, manufacturers, logistics firms, utilities. This category is where most analysts and data engineers on earth actually work, and it is almost never discussed, because a regional insurer does not run a developer marketing campaign.

They have decades of data, real problems, formal processes, and hiring that runs on qualifications and references. Slower, steadier, and far less competitive per opening than a known technology brand.

6. Agencies and consultancies

Small firms doing project work for other companies. Broad exposure, many stacks, fast learning, and constant context-switching. Good for building range early. Hard to go deep.

7. Public sector, universities, NGOs, research institutes

Census bureaus, health ministries, transport authorities, agricultural research bodies, international organisations. Lower pay, slow processes, genuinely interesting data, and unusually low competition because almost nobody applies. Also unusually good for a first portfolio-driven hire, because these organisations publish their data and you can build something useful before you ever apply.

8. Yourself


Freelance and contract. Covered in its own lesson. Not a fallback — a category with its own economics.

The ratio nobody prints

For every organisation that trains its own foundation models, there are thousands that merely run software and have accumulated data they cannot use. Categories 3, 5 and 7 hold most of the jobs and generate the least conversation.

Figure 1.2

Eight employer types, ordered by how many beginners they take



Categories three, five and seven hold most of the jobs and generate the least conversation. The famous box at the bottom employs a few thousand people worldwide, nearly all with doctorates.

If you are optimising for a first job, apply where the jobs are, not where the news is.

The "AI team" that is not one

A caution you will meet quickly. Many companies now advertise AI roles while having no model in production and no data warehouse. The actual job is spreadsheets, one integration with a hosted model API, and a dashboard. This is not fraud; it is where those companies genuinely are. But go in knowing it, because it decides what you learn in your first year. The interview lesson on questions to ask exists mostly for this.

Finding these employers without paying anyone

The paid tools — LinkedIn Recruiter, Sales Navigator, Apollo, paid job aggregators — sell you a filtered list of companies. You can assemble the same list free:

- LinkedIn's free search, filtered by industry and employee count, in your city.
- Company career pages directly. Most enterprise roles are posted there and syndicated late or never.
- Your national employment portal, industry association member lists and chamber-of-commerce directories.
- Government tender and procurement portals, which name every firm doing data work for the state.
- Wellfound and Y Combinator's job board for startups; Kaggle and Hugging Face job boards for applied roles.

An afternoon spent building a list of forty local employers in categories 3, 5 and 7 is worth more than any subscription.

BEFORE YOU MOVE ON

Ayo lives in Lagos, has one deployed project and no degree. He spends his search on the careers pages of five frontier AI labs and three global technology brands. Ninety applications, no replies. What is the most useful diagnosis?

- A. He needs a stronger portfolio before those companies will respond.
- B. His applications are being rejected by resume-screening software he could beat with better keywords.
- C. He is applying almost entirely into the two employer categories with the fewest openings and the highest bar, while the categories that hire most people at his level — services firms, enterprises with internal data teams, and public bodies — are absent from his list.
- D. Lagos is the wrong market and he should focus only on remote roles with foreign companies.

The answer and why: p. 470

Lesson 3 · 10 min

Geography, time zones, and the office you have never heard of

THE ONE THING TO KEEP

Where you sit decides which employers can legally and practically hire you, and time-zone overlap filters candidates harder than any skill requirement in the advert.

Four ways to be paid by a company that is not in your city

The employer categories tell you who hires. This lesson tells you where the desk is, because that single fact changes your pay, your hours and how hard the door is to open.

CASE STUDY · MODULE 1

Two offers, one title, and the question about the dashboard

An illustrative case, built from documented patterns.

Sneha, twenty-three, a commerce graduate in Nashik with eight months of evening SQL practice, holding two offers in the same week.

Offer A was from a fourteen-person edtech company in Pune. The title was **AI Engineer**. The money was ₹6,50,000 CTC. The founder had spent twenty minutes of the interview describing what they would build with language models next year.

Offer B was from the Pune office of a Dutch insurer — a global capability centre with about nine hundred people, a category Sneha had not known existed six weeks earlier. The title was **MIS Executive**. The money was ₹5,80,000 CTC.

Her mother had one question about each, and it was the same question: which one has AI in it.

What the adverts said when read backwards

Sneha had learned to read the tools list first. The startup's requirements named a hosted model API, Google Sheets, Zapier and "prompt libraries", and the responsibilities section ran to fourteen bullets covering dashboards, marketing support, customer research and "AI initiatives". The insurer's advert named SQL, dbt, Snowflake and Airflow, listed four responsibilities that described one coherent activity, gave the manager's name, and described the first ninety days.

She also noticed something she nearly did not mention to anybody: the startup's advert had been reposted every month since April.

The question she asked in both interviews

Where does the number on your main dashboard come from, end to end?

The founder said it depended who you asked, and laughed, and said that was exactly the problem she would be hired to fix. The insurer's team lead answered in four sentences: the policy administration system lands in Snowflake by 04:00, twenty-six dbt models run against it, forty-one tests run after that, and if a test fails the dashboard shows yesterday's data with a banner rather than a wrong number.

One of those answers describes rung two of the maturity ladder. The other describes rung three or four.

The cost on both sides

Taking the startup meant twelve per cent more money, a title that would look serious on a CV, and touching a model API in her first month. It also meant being the only technical person on data, with nobody to review her work, in a company where the first year would be spreadsheets and one integration — and where the advert being reposted monthly suggested she would not be the first person to discover that.

Taking the insurer meant a title that sounded like a clerk's, no model anywhere near her desk for at least a year, and the slow, procedural culture of a large European employer. It also meant a named manager, a team of six, a reviewer for everything she wrote, and a system she could learn deeply.

The startup wanted an answer within twenty-four hours. The insurer had given her a week.

What she was actually choosing between

Not two jobs. Two first years. One of them would produce a person who had spent twelve months doing four things badly with nobody to correct her, holding a title that would cause a large employer to down-level her at the next application anyway. The other would produce a person who could say, in an interview, exactly where a number came from and what happens when the pipeline fails.

What she could not know was whether the startup would be the exception — the small, chaotic place where somebody learns five years of work in two.

YOUR ANSWERS

1. Sneha's mother was right that the word AI matters to employers. Where in the process does a title actually help, and where does it stop helping?

2. The founder said the vague answer about the dashboard was the problem she was being hired to solve. Why is that answer not enough on its own?

3. What would have changed the decision in the startup's favour?

STOP HERE

Write your answers before you turn the page. How it played out starts on p. 54.

CASE STUDY · MODULE 1

How it played out

She took the insurer, and asked for the offer in writing before telling the startup. Her first eight months were SQL, dbt models and a rota for checking test failures — no models, no AI in the work at all. In month eleven she moved internally onto the data platform team after answering a question in a meeting with a number and a source, which is the only interview that move required. The startup's advert was reposted twice more that year. Two years later she interviewed at a product company and was mapped to their mid-level band on scope; the interviewer's only note about her CV was that she had a system she could describe end to end, which almost nobody applying at that level could.

The questions, discussed

1. Sneha's mother was right that the word AI matters to employers. Where in the process does a title actually help, and where does it stop helping?

A title helps at the top of a funnel — it gets a CV opened and it matches the words a recruiter is filtering on. It stops helping the moment somebody maps her to their own internal level, because large employers map on scope rather than on the title she held. An inflated title also invites interviews for jobs whose scope she cannot hold, and those interviews are miserable and consume the weeks she needed for the ones she could pass.

2. The founder said the vague answer about the dashboard was the problem she was being hired to solve. Why is that answer not enough on its own?

It is a truthful description of the job and it is also a description of rung one or two of the maturity ladder with no senior person on it. The work is real and worth doing, but the advert sold model work and the year would be extraction and reconciliation. The question is not whether the job is legitimate; it is whether she would end the year able to explain what she had learned, and with somebody having reviewed any of it.

3. What would have changed the decision in the startup's favour?

A senior technical person she would report to, an honest advert that described the plumbing rather than the models, and an advert that had not been reposted monthly for five months. Any one of those changes the risk materially; the first changes it most, because the defining cost of a junior role with no reviewer is a year of habits nobody corrects.

MODULE 1 TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record. The answers, and the reason behind each, are in the key on p. 468. If two go wrong, re-read the module before the exam.

- 1 You have forty applications to spend this month. Why do enterprise internal data teams and IT services firms usually give a better return per application than a well-known technology brand?
 - A. Far fewer people apply per opening, because those firms run no developer marketing.
 - B. They pay above the market, which makes candidates self-select away from them.
 - C. They hire purely on portfolios, so qualifications are not filtered at any stage.
 - D. They do not use tracking systems, so a human being reads every application that arrives.
- 2 A San Francisco team requires overlap from 09:00 to 13:00 Pacific time. For which candidate is that a night job rather than an inconvenience?
 - A. São Paulo, where it falls in the early afternoon
 - B. Lagos, where it falls in the evening
 - C. Bengaluru, where it falls between 21:30 and 01:30
 - D. Warsaw, where it falls late in the evening and working-time rules may apply
- 3 An advert is titled "AI Engineer". The requirements name dbt, Snowflake, Airflow and strong SQL, with Python mentioned once at the end. What job is being advertised?
 - A. A prompt and evaluation role, because AI is in the title
 - B. A research position, since Airflow is used to schedule training runs
 - C. An analytics role, because SQL appears before the other tools
 - D. A data engineering job, whatever the header says

- 4 Fully remote junior roles for foreign employers became scarcer while remote senior roles did not. What is the mechanism?
- A. Junior work needs supervision that cannot happen over a video call.
 - B. An employer-of-record fee is a fixed per-head cost, so it weighs far more against a junior salary.
 - C. Immigration rules were tightened for remote workers in most countries.
 - D. Junior candidates apply in such overwhelming numbers that employers stopped advertising the roles publicly.
- 5 Why does two hours every week beat twelve hours once a month, even though the monthly total is similar?
- A. Facts are retained better when they are revisited at spaced intervals rather than in one long block.
 - B. Twelve-hour sessions exceed the point at which concentration is productive.
 - C. A gap costs a whole session, because you must rediscover your own reasoning before adding a line.
 - D. Employers can see commit dates and prefer a steady contribution history.
- 6 Two people have the same tools and the same five years. One is levelled senior and one is levelled mid. What distinguishes them?
- A. The senior one is given vague areas and defines the problem others then agree to.
 - B. The senior one writes code that is correct more often and needs less review.
 - C. The senior one knows more algorithms and more of the current tooling.
 - D. The senior one has managed people, which is the step above individual contribution everywhere.

BEFORE YOU MOVE ON

A nightly pipeline writes only 40% of the expected rows because a source API rate-limited it, but the job exits successfully and the dashboard updates. Why is this worse than the job crashing outright?

- A. Because a crash would have triggered the retry logic, and retries always recover rate-limit failures.
- B. Because incomplete data that looks complete is acted on, and the error is only discovered later when trust is already spent — whereas a crash is noticed immediately and blocks the bad numbers from being used.
- C. Because partial writes corrupt the table's schema and require a full rebuild.
- D. Because the pipeline used more compute than budgeted for a partial result.

The answer and why: p. 477

Lesson 14 · 10 min

The machine learning engineer

THE ONE THING TO KEEP

Most of a machine learning engineer's work is the system around the model, and the defining bug is a feature that exists in training data but not at prediction time.


The model is the small part

The title suggests a person who trains models. The job is mostly building and operating the system around a model: getting features to it reliably, serving predictions inside a latency budget, noticing when quality decays, and being able to roll back.

A useful proportion, borne out by most teams that have measured it: of the effort in a deployed machine learning system, model training is somewhere between five and fifteen per cent. The rest is data, plumbing, evaluation, monitoring, and the arguments about what should be predicted at all.

Figure 2.1

Where the effort goes in a deployed model



Training is somewhere between five and fifteen per cent of the work, and it is the part every course teaches. The rest is the system around the model, and that is the job.

The week

You are asked to reduce failed deliveries. Week one is not modelling; it is establishing what "failed" means in the data and discovering that the field was only populated properly after last April. Week two you build a baseline — predict the customer's historical failure rate — and it is 80% as good as anything you will build later, which is worth knowing before you spend a month. Week three you add features and gain four points. Week four you find one of those features is not available at prediction time in production, only in the training table, and you delete it and lose the four points.

That last one has a name and it is the defining bug of the field.

EXAMINATION

Working in AI — final examination

This paper covers the whole course:

- the employer types and roles that exist and how to read an advert for the job it really is
- what each role does in an ordinary week
- the evidence a reviewer looks for and how to produce it
- the routes that put a person in front of your work
- using a model in your own search without destroying the thing being judged
- every stage of a hiring process from the recruiter screen to the signed offer
- what is actually in an offer, and the contract and equity terms underneath it
- freelancing priced as a business
- the first ninety days and the ten years after

65 questions, the whole pool. The site draws 25 for a sitting, pass mark 70%.
Work any 25 under your own conditions. Answers from page 526.

1 Module 1 · The terrain: who hires, for what, and where you fit

Across data analysis, data engineering, ML engineering and AI engineering, what does the largest share of the working day have in common?

- A. Something is wrong with the data and you have to find out why
- B. Communicating results to stakeholders who lack technical background
- C. Choosing between competing tools and frameworks for a given task
- D. Writing and reviewing production code under version control

2 Module 1 · The terrain: who hires, for what, and where you fit

What is the durable value of a year of annotation or model-evaluation work on a platform like Outlier, Appen or Toloka?

- A. A promotion track from labelling into research engineering at the same company
- B. Failure-mode intuition and the vocabulary of data quality, taken elsewhere
- C. A reference from a well-known AI company that recruiters recognise
- D. Enough accumulated income to fund a period of full-time study afterwards

3 Module 1 · The terrain: who hires, for what, and where you fit

Which single interview question best reveals where a company sits on the data maturity ladder?

- A. How many people are on the data team, and how is that team structured?
- B. Which tools are you using for orchestration and transformation?
- C. What are the biggest data challenges the team is facing this year?
- D. Where does the number on your main dashboard come from, end to end?

4 Module 1 · The terrain: who hires, for what, and where you fit

A degree buys three separate things. Which of them is a genuine hard gate rather than a preference?

- A. Legal eligibility, where immigration rules name a recognised qualification
- B. The mathematics, which is difficult to acquire outside a university setting
- C. Credibility with hiring managers at product companies and start-ups
- D. Access to campus placement and graduate schemes at large employers

5 Module 1 · The terrain: who hires, for what, and where you fit

What is the most common way a promising year of part-time study produces nothing hireable?

- A. Choosing a lane that does not match the local job market's actual openings
- B. Spending too long on fundamentals before attempting anything practical
- C. Studying three lanes at thirty per cent each, because it feels prudent
- D. Building projects that are too small to demonstrate real technical ability

6 Module 1 · The terrain: who hires, for what, and where you fit

The same advert for a junior role has been reposted every month for six months. What does that most usefully tell you?

- A. The company is growing quickly and hiring several people into the same role
- B. The posting is being refreshed automatically to keep it near the top of listings
- C. The salary band is below market and candidates are declining offers
- D. Either people keep leaving the role, or the specification is impossible

Module 1 • The terrain: who hires, for what, and where you fit

The module starts on p. 11.

MODULE 1 TEST

Module 1 test, Q1, p. 56

You have forty applications to spend this month. Why do enterprise internal data teams and IT services firms usually give a better return per application...

Answer A: Far fewer people apply per opening, because those firms run no developer marketing.

Why: A regional insurer or a staffing firm does not run a developer marketing campaign, so almost nobody thinks of applying there. The competition per opening is a fraction of what a known brand attracts, and the work is often the same. Note that these employers usually do filter on qualifications and do use tracking systems — the advantage is the ratio of applicants to openings, not a friendlier process.

Module 1 test, Q2, p. 56

A San Francisco team requires overlap from 09:00 to 13:00 Pacific time. For which candidate is that a night job rather than an inconvenience?

Answer C: Bengaluru, where it falls between 21:30 and 01:30

Why: Bengaluru is UTC+5:30, so 09:00 Pacific is 21:30 local and the window ends at 01:30. That is not a late shift, it is a night shift, and people burn out of it inside two years. The arithmetic is why a South Asian candidate's realistic remote market is Europe and the Gulf far more often than the US west coast — a London team's 10:00 to 14:00 lands comfortably in both Bengaluru and Lagos.

Module 1 test, Q3, p. 56

An advert is titled "AI Engineer". The requirements name dbt, Snowflake, Airflow and strong SQL, with Python mentioned once at the end. What job is...

Answer D: A data engineering job, whatever the header says

Why: The tools list is the most honest section of any advert. Whoever wrote it may have been vague about the work, but they were not vague about the stack, because the stack is what the team argues about. Airflow, dbt and a warehouse is data engineering. Apply to it as one, and expect the interview to test pipelines rather than prompts.

Module 1 test, Q4, p. 57

Fully remote junior roles for foreign employers became scarcer while remote senior roles did not. What is the mechanism?

Answer B: An employer-of-record fee is a fixed per-head cost, so it weighs far more against a junior salary.

Why: Hiring where a company has no legal entity risks creating a taxable presence and inherits that country's employment law, so companies route people through an employer of record, which charges a few hundred dollars a month per person. That fee is trivial beside a senior salary and material beside a junior one. It is an arithmetic filter, not a judgement about whether juniors can work remotely.

Module 1 test, Q5, p. 57

Why does two hours every week beat twelve hours once a month, even though the monthly total is similar?

Answer C: A gap costs a whole session, because you must rediscover your own reasoning before adding a line.

Why: The cost of a gap is re-entry, not decay of the calendar. Memory of your own reasoning — why you chose that join, what you were halfway through fixing — fades much faster than memory of facts, so the first session back is spent rebuilding context rather than making progress. Two gaps a month and half the year goes on re-entry. Writing one NEXT line at the end of every session is the cheapest defence.

Module 1 test, Q6, p. 57

Two people have the same tools and the same five years. One is levelled senior and one is levelled mid. What distinguishes them?

Answer A: The senior one is given vague areas and defines the problem others then agree to.

Why: Levels are set by how much ambiguity you absorb and how far the consequences of your work travel. Only the junior level is about producing a correct result from a defined task; everything above is about deciding what should be produced and getting others to agree. This is why excellent technical people plateau at mid-level, and why management is a fork rather than a step — many companies pay both tracks equivalently for several levels.

THE LESSON CHECKS**Lesson 1, p. 15**

A 300-person logistics company in Jakarta hires its first AI person to build an assistant that answers staff questions about shipments. Three months in, where...

Answer B: Getting shipment data out of four systems that disagree with each other, and building a set of test questions to check the answers against.

Why: Most AI work is moving messy data and checking outputs, not training models.

A Fine-tuning a language model on the...: Assumes the hard part of AI work is training models, because that is what courses spend their weeks on.

C Comparing model providers and negotiating a...: Treats the job as vendor selection — a real task, but one that takes days, not months.

D Rewriting the prompt until the answers...: Mistakes the visible surface of the system for its main cost, because the prompt is the part you can watch changing.

Lesson 2, p. 20

Ayo lives in Lagos, has one deployed project and no degree. He spends his search on the careers pages of five frontier AI labs and...

Answer C: He is applying almost entirely into the two employer categories with the fewest openings and the highest bar, while the categories that hire most people at his level — services firms, enterprises with internal data teams, and public bodies — are absent from his list.

Why: Nearly every AI-adjacent job sits inside an ordinary organisation that uses software, not inside a company that builds models.

A He needs a stronger portfolio before...: Assumes the barrier is evidence quality, when the categories he chose filter mainly on credentials, referrals and fixed graduate intakes.

B His applications are being rejected by...: Treats a structural mismatch as a formatting problem, which is the most commonly sold and least useful explanation.

D Lagos is the wrong market and...: Skips over the large local employer categories in favour of the most globally competitive option available.

Lesson 3, p. 23

Priya is in Bengaluru with two years of analytics experience. She applies only to fully remote roles at US start-ups and gets nowhere. Which change...

Answer A: Target European and Gulf employers and the global capability centres of foreign companies already operating in India, because both give workable overlap hours and an existing legal entity to employ her through.

Why: Where you sit decides which employers can legally and practically hire you, and time-zone overlap filters candidates harder than any skill requirement in the advert.

B Improve her portfolio until it is...: Assumes the barrier is candidate quality when the binding constraints are a nine-and-a-half hour offset and the absence of a legal entity to employ her.

C Accept a night shift, because overlap...: Treats a genuinely damaging schedule as unavoidable when a different target market removes the problem entirely.

D Register as a freelancer so US...: Solves the entity problem but silently converts a salaried role into contract work with no notice, benefits or severance, which is a large trade presented as a formality.

USE IT IN YOUR WORK

Working in AI

The roles that exist, what they pay, and how people get in without a famous degree.

WHAT IT TEACHES

A plain-spoken guide to building a career in and around AI: the roles that actually exist and what they look like on an ordinary Tuesday, where a degree matters and where it does not, how to build a portfolio with no job history, what to do when the courses run...

WHAT IS INSIDE

Nine modules and 85 lessons, 33 figures drawn from data, nine case studies, nine module tests, a 65-question examination, and every answer with the reason behind it.

LICENCE

CC BY-NC-SA 4.0. Copy, print, share and adapt it for any non-commercial purpose, with credit, under the same licence. There is no paid edition.

THE COURSE

Free to read, with the lessons, the films and the examination, at addaly.com/learn/ai-careers

Addaly

First edition, September 2026 · build 62a26075



Scan to open the course