USE IT IN YOUR WORK

AI for a Small Business

Two or three tasks, honestly measured. Not a new department.

First edition, September 2026

Addaly

Series editor: Dr Mustafa Mun

Draft, open for academic review

USE IT IN YOUR WORK

AI for a Small Business

Two or three tasks, honestly measured. Not a new department.

Series editor: Dr Mustafa Mun

Addaly

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

AI for a Small Business

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). AI for a Small Business. Addaly. addaly.com/learn/ai-for-small-business.*

This book is the printed form of the course of the same name at addaly.com/learn/ai-for-small-business. 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 · 88 lessons · 29 figures · 9 case studies · 69,888 words · about 13 hours

Brief contents

1

MODULE 1

Working out where AI belongs in your business

Produce a one-page map of your own week that names two candidate tasks, states the hours each takes today from a real five-day count, prices the checking those tasks will need, and explains from the model's mechanism which of your tasks it will fail at.

Lesson 1 · 7 min

The Two Tasks Worth Doing

THE ONE THING TO KEEP

Start from where your hours actually go — frequent, text-shaped, cheap-to-check work — not from what the tool can do.

The wrong question

The usual question is "what can AI do for my business?" It gives you a list of forty possibilities and you act on none of them.

Ask a different one. Where did last week actually go?

Keep a rough tally for five working days. Not a time sheet — a note on your phone each time you switch task. Owners are usually surprised by two things: how much of the week is made of words (reading messages, writing replies, explaining the same thing to the eleventh person), and how little of it is the work they think they do.

Three filters

Put every item on that list through three questions.

Do you do it more than twice a week? Savings are per-instance, so frequency is the whole game. Ten minutes saved on something daily is about 40 hours a year. Two hours saved on something you do every December is two hours.

Is it made of words? These tools read and write text. They sort, summarise, draft, translate, extract, rephrase. They do not lift boxes, chase a supplier who is ignoring you, or know that Mrs Okafor always wants her order left with the neighbour.

If it came back wrong, would you notice cheaply? You will be checking the output. If checking takes as long as doing, you saved nothing. If a mistake is silent and expensive — a wrong tax figure, a wrong dimension on a fabrication order — the checking has to be so careful that the tool stops paying.

A task that passes all three is a real candidate. A task that fails one is not, however good the demo looked.

Figure 1.1

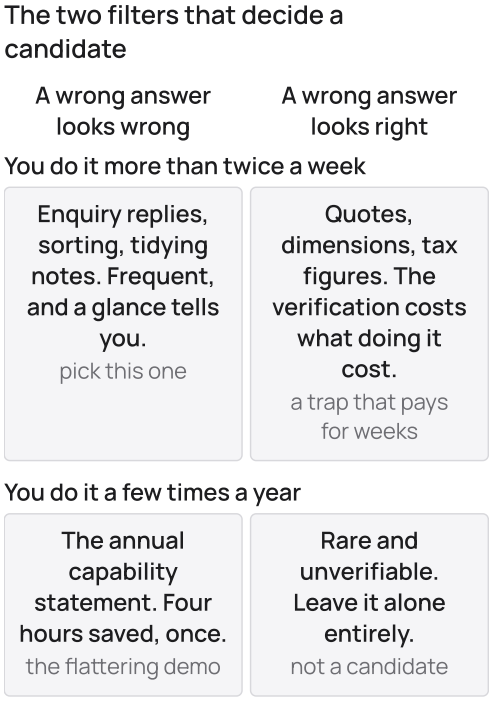

Frequency multiplies the saving and the cost of checking divides it. A task that fails on either axis is not a candidate, however good the demo looked, and no amount of better prompting moves it.

What this looks like

Three shapes of business:

- A tailor in Lagos gets 30–40 WhatsApp messages a day. Twenty-five of them are "how much for a two-piece?", "how long?", "do you deliver to Ikeja?" Frequent, text, obvious when wrong. Strong candidate.
- A physiotherapy clinic in Manila writes a short visit summary after each appointment. Frequent and text-shaped — but it is a clinical record, where a wrong detail is silent and serious. Candidate with a hard rule attached: the therapist writes the clinical content, the tool tidies the language.
- A three-person agency in Bogotá writes six proposals a month. Frequent enough, text-shaped, easy to check. Strong candidate — as long as the price and the dates never come out of the model. Lesson three covers that.

The arithmetic, before you spend anything

Do this on paper.

Hours the task takes now, per week. Hours you honestly think it will take with help — including the editing, which is where new users always undercount. Multiply the difference by 48 weeks. Then price your own hour: not what you charge a client, but what you would pay someone to take that hour off you.

A worked case. Answering enquiries takes 5 hours a week and drops to 3. That is 2 hours a week, 96 hours a year. If your hour is worth ₹600, that is roughly ₹57,000 of time against a tool costing perhaps ₹24,000 a year. Clear yes. Run the same sum on "writing my annual capability statement" and you get 4 hours saved, once. Clear no — and it was the more impressive demo.

What to leave alone

Deliberately ignore:

- Anything you do a few times a year, however painful.
- Anything a cheaper, dumber tool already fixes. A saved reply in WhatsApp Business costs nothing and answers "what time do you open" perfectly, forever. Lesson eight comes back to this.
- Anything where you cannot check the answer. If you do not know enough about your local tax rules to spot a wrong answer, a confident wrong answer is worse than no answer.

- Anything that is really a people problem. No tool fixes a supplier who does not call back.

Pick two

Two tasks. Three at the very most. One person, one tool, for the first month.

That is not caution for its own sake. It is that you only learn what these tools are like on *your* work by using one on your work, repeatedly, until you can predict what it will get wrong. That prediction is the asset. Eight half-tried tools buy you eight shallow impressions and no judgment at all.

BEFORE YOU MOVE ON

Priya runs a 12-table restaurant and has four jobs she would like off her plate. Which is the strongest first candidate?

A. Rewriting the descriptions for all 40 dishes on the menu, which she last touched two years ago

B. Answering the 15–20 daily messages about opening times, parking and whether there are vegan options

C. Setting next quarter's prices from her supplier costs and what nearby places charge

D. The annual tax return, which she dreads for three weeks beforehand

The answer and why: p. 497

Lesson 2 · 9 min

Fluent and wrong come from the same place

THE ONE THING TO KEEP

A language model predicts likely next words rather than retrieving facts, so its confidence is a property of the writing and never evidence about truth — which is why the facts must go in for the words to come out safely.

Two words at a time

You do not need the mathematics. You do need one mechanical fact, because everything useful and everything dangerous about these tools follows from it.

A language model reads what is in front of it and predicts the next fragment of text. Then it reads that, and predicts the next. A fragment is called a **token** — roughly three quarters of a word, so a 400-word email is about 530 tokens. It does this a few hundred times a second, and a paragraph appears.

That is the whole trick. There is no lookup, no fact store, no moment where it checks anything. It has read an enormous amount of writing and learned which words tend to follow which, in which contexts.

Now ask it what time your shop opens.

It has never seen your shop. But it has read tens of thousands of pages about shops, and on those pages the sentence "We are open" is very often followed by "from 9am to 6pm, Monday to Saturday." So that is what comes out — in your voice, formatted correctly, entirely invented.

This is not the model malfunctioning. It is the model doing precisely the one thing it does. People call these inventions hallucinations, which makes them sound like a rare glitch. They are the ordinary output of a machine whose job is plausibility.

Why it cannot warn you

The obvious question: why does it not say "I am unsure about that"?

Because the only number it has is the probability of the next token, and that is a statement about word choice, not about truth. When it writes "9am", the model is confident that "9am" is a likely word in that position. It has no separate signal that says "this fact is unverified." There is nothing inside it that knows the difference between recalling and constructing.

So the confidence in the writing is a property of the writing. A fabricated price arrives in exactly the same calm, well-punctuated sentence as a true one. There is no tell. Anyone who tells you they can spot invented output by how it reads is describing luck.

Same question, different answer

Try this once, deliberately. Ask the same question three times in three fresh chats. You will get three different replies.

That is because the model does not always take the single most likely next token; it samples from the likely ones. Some tools expose this as a "temperature" or "creativity" setting. Turned down, output is repetitive and safe. Turned up, it is varied and more prone to wandering.

The business consequence is the one people miss: **you cannot test this thing once.** A single good answer tells you almost nothing, because a different one was available. When you evaluate a prompt later in this course, you will run ten cases, not one.

What it has read, and when it stopped

Every model has a training cutoff — a date after which it has read nothing. Ask about a tax change from three months ago and you may get a fluent description of the old rule, or a fluent description of a rule that never existed.

Some tools now attach extras: a web search, a calculator, your uploaded files. When those are switched on the picture improves, because the model is reading a real page instead of reaching for a memory. It is worth knowing

whether the tool you are using has search on or off, because the same question gives materially different answers either way, and the interface rarely makes it obvious.

What this predicts about your work

The mechanism sorts your tasks for you.

It will be strong wherever the **shape** of the answer is the hard part and the facts come from you: turning your notes into a scope section, rewriting an angry reply calmly, sorting 300 transaction descriptions into categories, producing forty product descriptions in the pattern of five you wrote yourself.

It will be weak wherever the **content** has to be specific, private, recent, or numeric: your prices, your lead times, this month's regulations, the total of a column, what your supplier said on the phone.

A reasonable rule that follows directly: **the facts go in, the words come out.** If you find yourself hoping a fact will come out, stop. You have just asked a plausibility machine for a matter of record.

The useful reframe

Stop thinking of it as something that knows things. Think of it as an extremely well-read assistant with no memory of your business, no access to your files, no ability to check anything, and an unshakeable reluctance to admit ignorance.

You would still hire that person. You would just never let them answer a customer unsupervised, and you would give them your price list first.

BEFORE YOU MOVE ON

You ask a model, with no documents attached, what your workshop's delivery lead time is. It answers "typically 7 to 10 working days" in a confident, well-formed sentence. What has actually happened?

A. It produced the continuation that is most plausible after that question given everything it has read about workshops, which is unrelated to your actual lead time.

B. It has retrieved an industry average from its training data, so the figure is a reasonable starting estimate for your business.

C. It guessed, and the confident phrasing means the guess sits near the top of its probability range, so the figure is more likely to be right than wrong and can be treated as a working estimate.

D. The model has malfunctioned and needs a clearer prompt specifying that it should only answer from verified sources.

The answer and why: p. 498

Lesson 3 · 9 min

The checking tax nobody counts

THE ONE THING TO KEEP

The saving from a drafting tool is the drafting time minus the checking time, and a task whose errors are silent rather than loud can cost more to verify than it ever cost to do.

The saving is not the drafting time

A drafted reply appears in four seconds. It used to take you four minutes. The obvious conclusion is that you saved four minutes.

CASE STUDY · MODULE 1

Fourteen hours, and the reason not to take them

An illustrative case, built from documented patterns.

Anand Stationers, a school-uniform and book shop in Patna with four staff, thirty-one school booklists pinned to a board, and an admission season that runs for six weeks.

Every June, Anand Stationers answers the same question about four hundred times: what does my child need for class six at St Karen's, and what will it cost.

The shop's WhatsApp number carries most of it. A parent sends a school name and a class. Somebody behind the counter finds the right booklist on the board, types out the titles, adds the uniform items, works out a total and sends it back. It takes between six and eleven minutes depending on how legible the booklist is and how many follow-up questions arrive.

Sumit Anand, who runs the shop with his father, spent five days counting properly rather than remembering. The tally came to 212 booklist enquiries in a single ordinary week of June, at an average of just under eight minutes each. That is twenty-eight hours, spread across four people who also have a shop floor to run. In the same week the shop wrote eleven credit-account reminder letters, which took forty minutes in total and which everybody hated.

The hated task was the one Sumit had planned to hand over. The tally moved it to the bottom of the list in an afternoon, which is what a tally is for.

So: the booklists. Frequent, obviously. Made of words, obviously. The third filter is where it stopped being obvious.

If a booklist reply comes back wrong, does anybody notice cheaply? A wrong greeting is visible in a second. A wrong title is not. A parent who is told their daughter needs the 2024 edition of a grammar book when the school has moved to the 2026 one buys the wrong book, discovers it in the second week of term, and comes back angry at a shop that had one job. At about two

thousand four hundred rupees a kit, a run of thirty wrong replies is a real number, and it is the kind of number that also costs the shop a family for five years.

Sumit's father made the objection plainly. The whole shop, he said, is thirty-one lists that change every year and nobody else in Patna keeps properly. If the machine starts guessing at them, we are selling the one thing we have.

The other side of the argument was not hypothetical either. Twenty-eight hours a week is one person for most of June. In the previous season the shop had lost custom simply by being slow — parents who sent a message at nine in the evening and got a reply at eleven the next morning, by which time they had gone to the shop opposite the bus stand. Sumit could name four.

So the choice was between leaving twenty-eight hours on the table for a second year, and putting a plausibility machine in front of the one record his customers trusted. Both sides had a cost and neither was small.

What broke the deadlock was drawing the task out as boxes rather than treating it as one job. Reading the message, deciding which school and class it refers to, finding the right list, typing it out, pricing it, and answering the follow-up. Six boxes, not one. Only two of them were drafting.

The box that actually consumed the minutes was the fourth, typing out between eleven and twenty-four lines from a photocopy pinned to a board. The box that caused the arguments was the second, where a parent writes "Karen's" and means one of two schools in Patna with similar names, or writes "class 6" when the school's list distinguishes between the two sections. Getting that box wrong sends a correct list to the wrong family, which is a failure the shop had been making about twice a week already, by hand, at four in the afternoon in the fourth week of June.

That changed the shape of the question. Sumit had been asking whether a model could be trusted to write a booklist reply. The map asked something narrower and answerable: whether it could be trusted to choose between thirty-one named documents, and to copy one out without changing it.

YOUR ANSWERS

1. The task Sumit wanted to hand over and the task worth handing over were different. What made the difference, and why could he not have known it without the tally?

2. The shop ended up using the model for classification and retrieval rather than for writing the reply. Why is that arrangement safer than a well-grounded drafting prompt would have been?

3. Sumit expected twenty-eight hours and got about sixteen. Why is that gap worth recording rather than rounding away?

STOP HERE

Write your answers before you turn the page. How it played out starts on p. 50.

CASE STUDY · MODULE 1

How it played out

The shop kept the booklists out of the model's mouth entirely. All thirty-one lists were scanned as dated PDFs into a folder, one per school per class, superseded versions deleted rather than renamed. The assistant's job was narrowed to two things: decide which school and class a message refers to, choosing UNSURE where two schools were plausible, and then quote the matching list verbatim with the file name and date attached. Anything it could not match returned NOT IN SOURCES rather than a list. Pricing stayed in the shop's own spreadsheet, because a total is arithmetic. Over the first fortnight the average enquiry fell from just under eight minutes to about three and a half, including the check, which came to roughly sixteen hours saved in a June week rather than the twenty-eight Sumit had hoped for — the difference being the checking he had not counted. The classification step turned out to be the surprise: it was wrong on nine of the first two hundred, and eight of those nine were the two Karen's, which UNSURE caught. The credit-account letters were never automated, on the grounds that forty minutes a year is forty minutes. The following January, when four schools changed their lists, the folder was updated on the same afternoon as the board, because Sumit had put it on the same checklist.

The questions, discussed

1. The task Sumit wanted to hand over and the task worth handing over were different. What made the difference, and why could he not have known it without the tally?

Frequency, and the direction in which memory fails. The credit-account letters were remembered as a burden because they are unpleasant, and unpleasant work is remembered by how it felt rather than by how long it took — forty minutes a year. The booklist replies were two hundred short interruptions, which register as nothing at all because no single one of them feels like a task. Memory inflates the hated and erases the frequent, so an owner acting on recollection reliably attacks the job that annoys them rather than the job that costs them.

2. The shop ended up using the model for classification and retrieval rather than for writing the reply. Why is that arrangement safer than a well-grounded drafting prompt would have been?

Because it removes the chance to invent rather than reducing it. A classifier can only return a label from a list the shop wrote, and a verbatim quote from a named file is checkable against that file in seconds. A drafting prompt, however well grounded, can still produce a plausible title that is not on any list, and a wrong title reads exactly like a right one. The shop moved the task from one where a wrong answer is invisible to one where it is either a wrong label or a wrong file name, both of which are visible.

3. Sumit expected twenty-eight hours and got about sixteen. Why is that gap worth recording rather than rounding away?

Because the missing twelve hours are the checking tax, and it is the number that decides whether the next task is worth attempting. The saving from any drafting or retrieval step is the original time minus the time spent reading, verifying and correcting, and new users undercount the second term almost without exception. Writing down that the real figure was sixteen gives the shop an honest ratio to apply to the next candidate, and a defensible claim rather than a feeling.

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. 496. If two go wrong, re-read the module before the exam.

1 A tailoring unit answers about forty price enquiries a day and writes one annual insurance renewal summary that takes half a day. It puts the enquiries first. What makes that the right order?

A. Savings are per instance, so a daily task compounds and a yearly one does not

B. A model handles short messages more reliably than it handles long documents of any kind

C. Insurance paperwork is confidential, and confidential material should never be shown to a model

D. The renewal summary needs judgement about the business, whereas a price reply does not

2 Asked for the opening hours of a shop it has never seen, a model produces "from 9am to 6pm, Monday to Saturday" in a calm, correctly punctuated sentence. What does that calm tell you about the answer?

A. That the model retrieved it from its training data and is therefore reporting a fact

B. That temperature has been set low, which suppresses the hedging it would otherwise add

C. Nothing, because the confidence is a property of the writing

D. That the model is more certain than usual, so the answer is more likely to be right

3 Timed end to end, a technical quote falls from eleven minutes to nine while a customer reply falls from four minutes to a minute and a half. Somebody proposes better prompting for the quote. Why will it not help much?

A. Because technical vocabulary sits outside what a general model has been trained on
B. Because the cost is in verification, and a prompt changes only the generation
C. Because quotes are longer, and output length is what drives how long a model takes to respond
D. Because the model needs a larger context window before it can handle specifications properly

4 A bookkeeping practice is choosing a door for a sorting task its accounting software already offers. Why is the built-in feature usually the cheapest route even when its headline price is not the lowest?

A. Because the model inside it has been tuned by the vendor specifically for bookkeeping work
B. Because bundled AI features are always priced below cost to keep customers on the platform
C. Because the data is already there, with no second vendor and no second copy
D. Because accounting vendors are forbidden by law from using your records to train models

5 An owner is certain that invoice chasing eats six hours a week. A five-day tally records two hours and twenty minutes. Which way does memory usually err, and why?

A. Downwards for all tasks, because people routinely forget the interruptions between them

B. Randomly, which is why a single week of tallying is not reliable enough to act upon

C. Upwards for everything, because owners want to justify the money they are about to spend

D. Upwards for disliked work, because unpleasant tasks are remembered by how they felt

6 A salon halves the time it spends drafting replies to enquiries that mostly ask about prices. Publishing the price list on its booking page then cuts those enquiries by seventy per cent. What did the first change miss?

A. That the drafting prompt had not been tested against ten real cases before it was used

B. That a faster reply increases enquiry volume, which cancels out the time saved on each one

C. That price enquiries are a classification problem rather than a drafting problem

D. That the step should have been questioned before it was made faster

Figure 2.1

The six sections of a one-page facts sheet

Prices
The base price, what pushes it up, the rush charge, the minimum order, delivery by area. Where a price is it depends, write down what it depends on: that sentence is what stops the invention.

Time
Opening hours, turnaround, lead times, how far ahead you book, and how long a quote stays valid.

Boundaries
Where you deliver and where you do not. Sizes you do not stock. Jobs you turn down. This section prevents more damage than any other, because a model asked about something outside your range will happily say yes.

Terms
Deposit, payment methods, your refund position in the words you actually use, cancellation, guarantee.

Voice
Three or four lines. Formal or casual, whether you use the customer's name, whether you use exclamation marks, what you never say.

The annoying questions
Three or four questions customers ask that irritate you, each with your real answer written out. These are the ones you answer worst when tired.

Thirty minutes, dated at the top, master copy in a file you own. Test it by taking your last twenty enquiries and asking whether the page answers each one; most first drafts answer about twelve, and the eight failures are the second draft.

Where to put it

Three options, same effect.

1. **Paste it at the top of every conversation.** Crude, free, works everywhere, and you will get bored of it by Wednesday.
2. **Standing instructions.** Most tools have somewhere the text lives permanently — custom instructions, a project, a saved assistant, a system prompt. Set it once and every conversation starts with your facts already in place.
3. **As a file.** Upload the page as a document and refer to it. Useful when it grows past a page, and it makes the dating visible.

Option two is right for most people. Whichever you choose, keep the master copy in a file you own, not only inside a vendor's product. You will need it again when you change tools.

Date it, and diarise it

Put the date at the top. Then put a repeating reminder in your diary — quarterly is about right.

The reason is a failure mode that is worse than having no facts sheet at all. A stale sheet is confidently wrong. Six months after a price rise, a page that says 7,400 will produce a reply quoting 7,400, in your voice, with none of the hesitancy a model shows when it is guessing. You have replaced an obvious risk with an invisible one.

When a price changes, the sheet changes the same day. Treat it like the till.

One sheet or several

Start with one. Split it when it goes past roughly two pages, or when two tasks need genuinely different material.

A useful split for a service business: a **customer sheet** (prices, hours, boundaries, voice) for anything a customer will read, and a **supplier and internal sheet** (cost prices, margins, which supplier is unreliable) for

EXAMINATION

AI for a Small Business: final examination

This paper covers the whole course:

- finding the two or three tasks worth handing over and pricing the checking they need
- writing, hardening and testing a prompt
- getting straight answers out of your own documents
- the enquiry-to-complaint path with a customer on the other end
- listings, images, advertising and price
- the paperwork of the business
- letting an automation act without losing control of it
- reading a year of your own numbers honestly
- what the whole thing costs, risks and looks like a year later

63 questions, the whole pool. The site draws 25 for a sitting, pass mark 70%. Work any 25 under your own conditions. Answers from page 555.

1 Module 1 · Working out where AI belongs in your business

You ask the same question in three fresh conversations and get three different replies. What follows for how you should evaluate a prompt?

A. That one good answer is not evidence and a prompt needs ten cases
B. That the tool is faulty and the vendor should be asked to investigate it
C. That temperature must be set to zero before any business use is safe
D. That the model is learning between conversations and improving as it goes

2 Module 1 · Working out where AI belongs in your business

Why does the numbers-and-promises pass come before the reading pass rather than after it?

A. Because numbers are at the top of most drafts, so it is the more efficient order
B. Because reading for tone first accepts the text, and the figures become furniture
C. Because tone is subjective and should be judged only once the facts are settled
D. Because a draft read twice in the same order produces diminishing attention each time

3 Module 1 · Working out where AI belongs in your business

On a process map, a step is marked R for rules. What does that marking imply about the tool it should get?

A. That it needs a model with a low temperature so that output does not vary
B. That it should be automated last, once the T steps have been proved
C. That it should be done by something that cannot vary
D. That it can be deleted, since anything fully rule-bound is usually unnecessary

4 Module 1 · Working out where AI belongs in your business

Your five-day tally records that two thirds of enquiries could have been handled by a competent stranger holding your price list. What is that fraction telling you?

A. The proportion of enquiries that are not worth answering personally at all
B. The accuracy you should expect from a classifier over the same messages
C. The share of your week that is already being duplicated by somebody else
D. Roughly the most any tool could take off you

5 Module 1 · Working out where AI belongs in your business

A model answers accurately about clause 14.3 of a ninety-page contract but misses two deadlines when asked to list them all. What explains the pattern?

A. Attention is diluted across a long input, and the middle is weakest
B. Lists are generated as a single output, and output length is capped by the tool
C. Deadlines are dates, and dates are tokenised inconsistently across a long document
D. Exhaustive questions are refused by default unless search is switched on

6 Module 1 · Working out where AI belongs in your business

Why is one person on one tool for one month better than four people trying four tools?

A. Because licence costs rise faster than the benefit across a small team
B. Because knowing what it gets wrong comes only from repetition
C. Because vendors give better support to accounts with a single named user on them
D. Because four tools produce inconsistent output, which customers notice quickly

Module 1 · Working out where AI belongs in your business

The module starts on p. 11.

MODULE 1 TEST

Module 1 test, Q1, p. 52
A tailoring unit answers about forty price enquiries a day and writes one annual insurance renewal summary that takes half a day. It puts the…
Answer A: Savings are per instance, so a daily task compounds and a yearly one does not
Why: Frequency is the whole of the arithmetic. Ten minutes saved on something done daily is about forty hours a year; four hours saved every December is four hours. The impressive task and the valuable task are rarely the same one, and frequency is what separates them.

Module 1 test, Q2, p. 52
Asked for the opening hours of a shop it has never seen, a model produces "from 9am to 6pm, Monday to Saturday" in a calm…
Answer C: Nothing, because the confidence is a property of the writing
Why: The only number the model has is the probability of the next fragment of text, which is a statement about word choice rather than about truth. An invented price arrives in the same well-punctuated sentence as a true one, and there is nothing inside the model that distinguishes recalling from constructing.

Module 1 test, Q3, p. 53
Timed end to end, a technical quote falls from eleven minutes to nine while a customer reply falls from four minutes to a minute and…
Answer B: Because the cost is in verification, and a prompt changes only the generation
Why: The eleven minutes did not become nine because the writing was poor. Every figure in a quote has to be traced back to a source before it can go out, and that costs roughly what it always cost. A prompt improves what is generated; it does not reduce what has to be checked.

Module 1 test, Q4, p. 53
A bookkeeping practice is choosing a door for a sorting task its accounting software already offers. Why is the built-in feature usually the cheapest route…
Answer C: Because the data is already there, with no second vendor and no second copy
Why: The advantage has nothing to do with the model. Stripping the data, pasting it, reading a second vendor's terms and keeping a second copy of your customer list somewhere you have not audited is most of the cost and all of the risk, and the built-in feature has already paid it.

Module 1 test, Q5, p. 54

An owner is certain that invoice chasing eats six hours a week. A five-day tally records two hours and twenty minutes. Which way does memory…

Answer D: Upwards for disliked work, because unpleasant tasks are remembered by how they felt

Why: Memory errs in a consistent direction, which is what makes four minutes a day of tallying worth it. Hated work inflates because it is remembered by feeling; frequent short interruptions vanish because no single one registers as a task. Together the two errors push owners toward the job that annoys them and away from the job that costs them.

Module 1 test, Q6, p. 54

A salon halves the time it spends drafting replies to enquiries that mostly ask about prices. Publishing the price list on its booking page then…

Answer D: That the step should have been questioned before it was made faster

Why: The right question about any box on a process map is not whether it can be drafted faster but why it exists at all. A step that disappears saves every one of its minutes and never needs checking, and making a badly designed step faster hides the problem for a while.

THE LESSON CHECKS

Lesson 1, p. 15

Priya runs a 12-table restaurant and has four jobs she would like off her plate. Which is the strongest first candidate?

Answer B: Answering the 15–20 daily messages about opening times, parking and whether there are vegan options

Why: Start from where your hours actually go — frequent, text-shaped, cheap-to-check work — not from what the tool can do.

A Rewriting the descriptions for all 40…: It is the most visible improvement, but work you do once every couple of years saves almost nothing across a year.

C Setting next quarter's prices from her…: It is the decision that moves the most money, which is exactly why an answer she cannot verify is expensive here.

D The annual tax return, which she…: It is the task she hates most, but it happens once a year and a mistake is both costly and slow to surface.

Lesson 2, p. 19

You ask a model, with no documents attached, what your workshop's delivery lead time is. It answers "typically 7 to 10 working days" in a...

Answer A: It produced the continuation that is most plausible after that question given everything it has read about workshops, which is unrelated to your actual lead time.

Why: A language model predicts likely next words rather than retrieving facts, so its confidence is a property of the writing and never evidence about truth — which is why the facts must go in for the words to come out safely.

B It has retrieved an industry average...: Treats generation as retrieval; there is no averaged figure stored anywhere, only a likely sequence of words, so the number carries no evidential weight at all.

C It guessed, and the confident phrasing...: Reads fluency as a confidence signal, but the probability governs word choice rather than truth, so smooth phrasing and invented content arrive together.

D The model has malfunctioned and needs...: Frames ordinary behaviour as a fault; a clearer instruction can reduce invention, but nothing in the prompt gives it access to a lead time it was never told.

Lesson 3, p. 23

A workshop uses AI to draft cutting lists from customer measurements. The prose is always clean, the checking feels quick, and it has run well...

Answer B: Cutting lists are numeric, so a transposed dimension reads exactly like a correct one, and six weeks of clean output is precisely what erodes the attention needed to catch it.

Why: The saving from a drafting tool is the drafting time minus the checking time, and a task whose errors are silent rather than loud can cost more to verify than it ever cost to do.

A Six weeks is too short a...: Sounds cautious but names no mechanism; a well-chosen task can be judged in two weeks, and duration alone is neither the risk nor the reassurance here.

C The model may have been updated...: A genuine risk covered later in the course, but it does not explain why this particular task is unsafe; the problem is silent numeric error, present from day one.

D The prose being clean shows the...: Misreads fluency as a tunable trade-off against accuracy, when clean prose and a wrong dimension are produced by the same process and neither causes the other.

USE IT IN YOUR WORK

AI for a Small Business

Two or three tasks, honestly measured. Not a new department.

WHAT IT TEACHES	A practical course for someone whose own money is at stake — a shop, a workshop, a clinic, a small agency, a nonprofit.
WHAT IS INSIDE	Nine modules and 88 lessons, 29 figures drawn from data, nine case studies, nine module tests, a 63-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-for-small-business**

Scan to open the course

Addaly

First edition, September 2026 · build 62a26075