YOUR QUICK TOUR

Trust is part of the product

Responsible AI is not a final warning label. It shapes the purpose, data, checks and people around the tool.

Meet Dani and the confident mistake 3

Start with purpose and risk 4

Protect private information 5

Check facts and sources 6

Look for unfair outcomes 7

Keep people accountable 8

Be transparent with customers 9

Use the traffic-light test 10

10 business ideas with first tests 11

Your seven-day safety sprint 13

Your responsible AI worksheet 14

YOUR FINISH LINE

One useful AI task with approved data, tested outputs, a named human owner and a clear fallback when the answer is uncertain.

01 / MEET YOUR GUIDE

Dani catches a confident mistake

Fictional example: Dani, her training business and all details are invented.



THE FAST DRAFT

Dani asks AI to draft a course brochure. The result sounds polished, but it invents an accreditation the business does not have.

THE HUMAN CHECK

She compares every claim with approved records, removes the false statement and adds a review step before anything is published.

THE LESSON

A confident answer can still be wrong. Speed is useful only when the process protects truth, people and trust.

02 / BEGIN BEFORE THE PROMPT

Ask what should be done

A responsible decision starts with the job, the possible harm and who carries the result.

NAME THE PURPOSE

Write the exact task and user. "Help draft a reply" is clearer than "use AI in customer service."

NAME WHO MAY BE AFFECTED

Include customers, staff, applicants, suppliers and people who may never see the AI tool.

NAME THE WORST REASONABLE FAILURE

Could the output expose private data, deny an opportunity, mislead a buyer, create a safety risk or damage a person?

COPY-READY PROMPT

"Before helping, list the people affected, sensitive data involved, three plausible harms, checks required, human approval point and safe fallback. Do not make the final decision."