

PRACTICAL AI RISK GUIDE

AI hallucination

A practical guide to the commercial risk of plausible but unverified AI output - and where a real source or human review must sit.

Confidence is not evidence.

Use AI for speed, then verify the details that carry consequence.

01**Name the risky detail.**

Facts, figures, quotes and commitments deserve a source.

02**Put review before use.**

A checkpoint matters before the output reaches a customer or decision.

03**Own the verification.**

A named person decides what is safe to publish or act on.

If a detail could create a promise, dispute or decision, check it first.

Use this guide →

MISSION & SCOPE

A polished answer is not the same as a verified answer.

AI can sound specific, confident and complete even when a detail is false. The risk begins when fluency is mistaken for evidence.

THIS GUIDE IS USEFUL IF

- AI supports research, summaries or customer copy.
- Specific claims travel into real work.
- The team needs a practical checking standard.

THIS GUIDE WILL NOT

- Estimate an error rate.
- Rank tools by accuracy.
- Treat every draft as equally risky.

How the risk travels

01

Generate

AI produces a plausible answer, summary or draft.

02

Decide

Someone checks the important claims - or assumes they are safe.

03

Use

The output reaches a customer, page, report or business decision.

Hallucination is a workflow risk, not simply a writing defect.

It is plausible but false, unsupported or invented information. Like a confident new hire who never says "I don't know", it can sound credible whether it is right or wrong.

BEFORE USE

Six details that deserve a verification decision.

Not every sentence needs the same scrutiny. Start with claims that can create commercial consequence.

Check what carries consequence.

Ask what would matter if this detail were wrong.

01

Facts

Can the statement be traced to a real source?

02

Figures

Is the number current, relevant and supported?

03

Quotes

Did the named person or document actually say it?

04

Commitments

Does the business really offer, guarantee or promise this?

05

Summaries

Was the interpretation checked against the original?

06

Ownership

Who decides whether the output is safe to use?

VERIFICATION BRANCH

The important branch happens before the output is used.

The same draft can be useful or dangerous depending on the path it takes.

LOWER CONSEQUENCE	HIGHER CONSEQUENCE
01 Light review Internal ideation or rough structure. Check obvious defects before reuse.	01 Source and owner required Customer promises, published claims or decisions need evidence and approval.
02 Verify when repeated A recurring detail can spread widely, so confirm it before standardising.	02 Stop unsupported use If the claim cannot be supported, remove it or return to the source.
THE CONTROL POINT Verification sits between generation and business use.	THE FAILURE PATH Confidence is treated as proof, so false detail travels unnoticed.

Need the human boundary? Read Where AI Fails →

COMMERCIAL EFFECT

Wrong detail becomes real cost when the business acts on it.

A rough idea is not the same risk as a false price, source or promise.

Source checked

Defensible claim

Important facts can be traced to something real before publication.

No source checked

Public correction

Specific detail reaches customers without evidence and must be unwound later.

Commitment reviewed

Promise controlled

Pricing, availability and guarantees are confirmed by a person.

Commitment assumed

Customer dispute

The draft invents terms the business never intended to offer.

Summary compared

Decision support

Important qualifiers survive because the original material is checked.

Summary replaces source

Hidden omission

A concise interpretation becomes settled fact without review.

Ask which details would need a source from a staff member.

Check now, before correction costs more.

FIVE SIGNS UNVERIFIED AI OUTPUT IS CREATING EXPOSURE

Problems to catch before the claim travels.

These are workflow failures: the tool drafts, but the business has not defined how truth is checked.

Critical

Sign 1

Facts, figures or quotes are published without a source

The output sounds specific enough that nobody confirms the evidence exists.

WHAT IT COSTS

Public corrections, damaged credibility and unsupported decisions.

Major

Sign 2

Customer-facing claims go out without review

AI drafts pricing, availability, guarantees or commitments unchecked.

WHAT IT COSTS

Confusion, disputes and promises the business never intended.

Major

Sign 3

No one owns the verification step

Everyone assumes somebody else will check the important details.

WHAT IT COSTS

Silent gaps in quality control and inconsistent team standards.

Moderate

Sign 4

Confidence is mistaken for accuracy

Specificity becomes a substitute for evidence.

WHAT IT COSTS

False information moves further because it sounds fluent and decisive.

Moderate

Sign 5

Summaries replace the original material

People rely on the summary where qualifiers or exceptions could matter.

WHAT IT COSTS

Important nuance disappears and later decisions inherit the wrong interpretation.

Specificity and accuracy are different properties.

Require a source or reviewer when money, reputation, customers or decisions are affected.

CORRECTION AND SELF-ASSESSMENT

Five questions before AI output travels further.

Common myth: “AI would not state something that specific if it were not true.” Better framing: detail can sound authoritative without being supported.

01

Which details could create a customer promise or commercial consequence?

These deserve the clearest evidence and approval step.

02

What real source would prove each important claim?

If nobody can name one, fluency may be standing in for evidence.

03

Who owns verification before publication or use?

“The team” is not ownership. Name the person or role.

04

Where are summaries replacing the original?

That is where omitted qualifiers can become invisible.

05

What changes after a false detail is found?

Correct the output and improve the workflow so the same error travels less often.

WORKING VERIFICATION STANDARD

A reusable check your team can apply before use.

Match the strength of review to the consequence of being wrong.

01 Identify the claim

Mark facts, figures, quotes, commitments and source-based summaries.

02 Rate consequence

Ask what happens if the detail is wrong or incomplete.

03 Name the source

Record the original page, document, record or person.

04 Assign review

Name who decides whether the content is safe to use.

05 Compare original

Check summaries and interpretations against the source.

06 Record learning

Use recurring mistakes to improve the process.

INSTEAD OF

“It sounds specific.”

No source, owner or consequence test.

Reusable verification record

TASK:

IMPORTANT CLAIM:

CONSEQUENCE IF WRONG:

SOURCE:

SOURCE CHECKED:

REVIEW OWNER:

APPROVED USE:

PROCESS CHANGE:

DISCUSS THE REVIEW →

Build the checkpoint.

Map claim, source and owner.

DUE DILIGENCE

Verify what carries consequence.

Three rules and two next steps for safer AI-assisted work.

01

Trace important claims

Require a real source for facts, figures, quotes and commitments.

02

Name the reviewer

Put a person or role between high-risk output and business use.

03

Check the original

Do not let a summary quietly replace the source where nuance matters.

Verification should follow consequence, not confidence.

If a provider cannot explain sources, review ownership and correction plainly, keep digging.

CONTINUE LEARNING

Your Data and AI

See what information your team may be sharing with AI tools.

[Read next →](#)

AI Foundations

Continue through the rest of this collection.

[Explore collection →](#)

Use AI for speed. Keep evidence in the loop.

Specific does not mean true. Put a source or named reviewer before customer promises, published claims and important decisions.

PRIMARY NEXT STEP

Your Data and AI

Understand what business and customer information is crossing into AI tools.

[Continue →](#)

AI Foundations

See every guide in this collection.

[Explore collection →](#)

All guides

Browse every topic and problem.

[View library →](#)

Need a safer verification workflow?

Bring one AI-assisted task where a wrong answer would matter. Map the claim, source, reviewer and correction path.

[Discuss the workflow →](#)