


AI Fog Detector

A practical field guide for spotting AI hype before it eats your budget, workflow, and last remaining adult judgement.



Published by Sienda Media 2026

Funny surface. Serious meaning. Concrete takeaway. Real-world evidence. Practical workaround.

Copyright

© 2026 Sienda Ltd. All rights reserved.

Published by Sienda Media for QOOBIX / Proteus.

QOOBIX is one of the trademarks of Sienda Ltd, a company registered in England & Wales under Company No. 08194971.

QOOBIX is a private market-intelligence application provisioned for selected business clients. It helps businesses create structured commercial intelligence reports for questions such as market entry, distributor discovery, partner discovery, competitor mapping, regional opportunity assessment, positioning, and action priorities.

www.qoobix.com

Proteus is the proprietary intelligence layer behind QOOBIX. It generates structured commercial intelligence from the client profile and the specific intelligence request.

Sienda Ltd, Third Floor, 207 Regent Street, London W1B 3HH, United Kingdom.

For all enquiries: media@siendaweblines.com

sienda.co.uk

Disclaimer

This guide contains satire, humour, cultural criticism, practical prompts, and references to AI, business, and organisational concepts. It is not legal advice, financial advice, technical implementation advice, cybersecurity advice, procurement advice, or a substitute for professional judgement. If an AI vendor begins speaking in glowing adjectives while

refusing to name the input, output, reviewer, risk, or cost, open a window and proceed with gloves.

A note from Proteus

Proteus is the narrative operating system behind QOOBIX: a shapeshifting myth-buster with a severe allergy to fashionable nonsense and a small drawer labelled 'phrases that died during due diligence'. In this guide, Proteus is not anti-AI. Proteus likes useful tools. Proteus dislikes expensive fog wearing AI perfume.

It is not here to tell you that machines are evil, that spreadsheets have feelings, or that the future should be faxed back to 1997 for safety. It is here to ask the questions the slide deck forgot to include: What task improves? What data is involved? Who checks the output? What breaks when the model is wrong? Why is nobody naming the process?

A chatbot placed near a business problem does not automatically become a strategy. Sometimes it becomes a chatbot placed near a business problem while everyone applauds the lighting.

Table of Contents

Introduction: The Fog Has Learned to Say “AI-Powered”	5
The Detector	7
1. “It is AI-powered.”	7
2. “AI will automate the process.”	9
3. “Prompting is the new strategy.”	11
4. “Agents will run the business.”	13
5. “The model will figure it out.”	15
6. “We have data, so we have intelligence.”	17
7. “AI will replace judgement.”	19
8. “The demo worked.”	21
9. “AI saves time automatically.”	23
10. “Everyone needs an AI strategy.”	25
What to Inspect Next	27
Sources, Further Reading, and Other Hype Detectors	28

Introduction: The Fog Has Learned to Say “AI-Powered”

AI can be useful. It can draft, classify, summarise, search, compare, assist, generate, route, flag, translate, simulate, and occasionally save a tired human from the spiritual death of rewriting the same paragraph twelve times. Used well, it can remove friction from clear work.

Used badly, it becomes a fog machine with a processor. It accelerates unclear processes, decorates weak products, multiplies bad data, hides responsibility behind a glowing interface, and gives ordinary confusion the thrilling feeling of being close to the future. The result is not intelligence. It is mist with a subscription plan.

AI fog — and theatre — begins when the word ‘AI’ arrives before the task. It says ‘automation’ before naming the workflow. It says ‘agent’ before naming the owner. It says ‘data-driven’ before checking whether the data is a compost heap with columns. It says ‘strategy’ before anyone has described the customer, the process, the risk, the review cost, or the decision that will actually change.

This guide is for people who want to use AI without being mugged by adjectives. It will help you distinguish useful AI from decorative AI, automated confusion, data mysticism, prompt worship, agent theatre, and chatbots wearing business model costumes. It will also make several pitch decks feel personally attacked, which is a public service.

What This Guide Is

This is a satirical field guide for inspecting AI claims before they eat your budget, workflow, and last remaining adult judgement. Each chapter takes a common hype claim, translates it into practical reality, identifies the risk, and gives you questions you can ask before buying, building, selling, deploying, or trusting an AI system.

It is designed for founders, consultants, professionals, managers, operators, creators, small business owners, and anyone who has heard the phrase ‘AI-powered transformation’ and felt the room temperature drop.

What This Guide Is Not

This is not anti-AI. It is anti-bullshit wearing AI perfume. It is not a coding manual, procurement playbook, cybersecurity checklist, AI ethics textbook, magic prompt book, or consultant deck with gradients and an endangered verb in the title.

It will not teach you to worship models, fear models, anthropomorphise models, or ask your chatbot to become a chief operating officer because someone on LinkedIn had a mystical moment with a workflow diagram.

How to Use This Guide

Each chapter gives you the hype claim, what it pretends to mean, what it often means in fog conditions, Proteus’s verdict, the real danger, the useful version, the questions to ask before spending money, and a practical next move.

Do not memorise the lines like spells. Use them as inspection tools. Calm questions are lethal to hype because they force the adjective to produce a process, an owner, a metric, and a receipt.

The Detector

1. “It is AI-powered.”

Claim

“It is AI-powered.”

What it pretends to mean

The product contains meaningful artificial intelligence that materially improves the outcome.

What it often means

The phrase “AI” has been attached somewhere near the product, possibly in the same building, and everyone is hoping the market will smell innovation before it asks what changed.

Proteus verdict

An adjective has been installed where a mechanism should be.

Why it is dangerous

“AI-powered” can hide the absence of task clarity. It may mean a useful model is embedded in a workflow, but it may also mean a normal feature has been given a futuristic hat. The danger is that buyers evaluate the aura instead of the result. Once the label starts doing the work of explanation, the system can escape ordinary scrutiny: input, output, reviewer, failure mode, cost, and accountability all disappear behind the glitter curtain.

The useful version

The useful version names the specific capability: classify invoices, summarise calls, identify duplicate records, draft first-pass responses, extract fields, route tickets, flag anomalies, or compare documents. “AI-powered” becomes meaningful only when it is tied to a defined task and a measurable improvement.

Ask this before spending money

- Which specific task is improved by AI?
- What input goes in, what output comes out, and who uses it?
- What metric proves the AI version is better than the non-AI version?
- What happens when the system is wrong?

Do this instead

Replace the adjective with a task sentence. Write: “This system uses AI to [task] for [user] so that [measurable improvement].” If the sentence collapses, do not buy the hat.

Field note

Composite field note: A company advertised an AI-powered customer platform. Under inspection, the AI feature was a generic chatbot that answered three support questions and then redirected users to the same help page they were already ignoring. The platform was not powered by AI. It was decorated with it.

Takeaway

AI-powered is not a claim. It is a doorbell. Make the mechanism answer.

2. “AI will automate the process.”

Claim

“AI will automate the process.”

What it pretends to mean

A messy workflow will become smooth, fast, cheap, and possibly fragrant because a model has been introduced.

What it often means

Nobody has mapped the current process, so automation is being invited to multiply the confusion at machine speed.

Proteus verdict

Automation is not a priest. It does not bless the mess before multiplying it.

Why it is dangerous

Automation can remove repeated work. It can also industrialise bad work. If a process has unclear inputs, hidden exceptions, conflicting owners, bad data, weak review, or political ambiguity, AI does not fix those problems by entering the room. It may simply execute them faster, make them harder to inspect, and generate new failure modes that look impressive until invoices, customers, or regulators become involved.

The useful version

The useful version begins with workflow design. Automation is valid when the process is understood, the task is repeatable enough, exceptions are known, review is proportionate, and the cost of error is acceptable.

Ask this before spending money

- What are the steps of the current workflow in plain human language?
- Which step is AI automating, assisting, or merely decorating?
- What exceptions require human judgement?
- Who reviews the output and what is the cost of review?

Do this instead

Before adding AI, write the workflow as task, input, output, decision, reviewer, risk, and stop condition. If the workflow cannot survive this sentence, do not automate it yet.

Field note

Fictionalised field note: A firm automated invoice review before standardising invoice formats. The model became extremely efficient at misunderstanding twelve kinds of inconsistency. Finance called it innovation. Suppliers called it Thursday with extra emails.

Takeaway

If the human workflow is fog, AI will industrialise the fog.

3. “Prompting is the new strategy.”

Claim

“Prompting is the new strategy.”

What it pretends to mean

The organisation can bypass planning, positioning, domain knowledge, process design, and decision-making by writing clever instructions into a model.

What it often means

Someone has discovered that asking a chatbot confidently for a strategy produces confident-looking paragraphs, which are now being mistaken for strategic thinking.

Proteus verdict

A prompt is not a strategy. It is a steering wheel drawn on a napkin unless connected to an actual vehicle.

Why it is dangerous

Prompts can be useful for exploring options, drafting, transforming formats, and generating first-pass material. But prompt worship treats language as if it were execution. It ignores context, constraints, customers, incentives, distribution, risk, and trade-offs. A beautifully prompted nonsense plan remains nonsense, only now it has subheadings and an unjustified sense of momentum.

The useful version

The useful version treats prompting as an interface, not a substitute for judgement. Good prompts support a defined task and are evaluated against domain knowledge, evidence, and outcomes.

Ask this before spending money

- What decision will this prompt support?
- What context must the model know?
- What expertise is needed to judge the answer?
- What would make the output unacceptable?

Do this instead

Create a prompt only after writing the decision it supports and the criteria for accepting the output. If you cannot judge the answer, the model is not assisting you; it is wearing your authority like a borrowed coat.

Field note

Composite field note: A consultant generated a market-entry strategy with a prompt and presented it as insight. The output looked polished, missed the legal constraint, invented a distribution channel, and recommended a partnership with a competitor that had ceased trading. The headings were excellent.

Takeaway

Prompting can improve work. It cannot replace knowing what work is.

4. “Agents will run the business.”

Claim

“Agents will run the business.”

What it pretends to mean

Autonomous AI agents will plan, execute, coordinate, optimise, and generally behave like a tireless operations team that does not require holidays or awkward conversations about salary.

What it often means

The business has unclear processes, but now wants unclear software entities to run them while everyone says “agentic” with the expression of someone handling sacred fire.

Proteus verdict

A model is not a manager. It should not be left alone with the steering wheel and a quarterly budget.

Why it is dangerous

Agent systems can be useful in constrained, supervised environments. They can also create hidden chains of action: calling tools, moving data, making decisions, sending messages, triggering workflows, and compounding errors before a human notices. The more autonomy a system has, the more important governance becomes. Otherwise you have not built an assistant. You have released a highly confident intern into procurement with API access.

The useful version

The useful version limits scope, defines permissions, logs actions, sets stop conditions, requires approval for high-risk moves, and assigns a human owner who is awake and reachable.

Ask this before spending money

- What actions can the agent take without approval?
- What tools, files, accounts, and data can it access?
- What are the stop conditions?
- Who is accountable for agent actions?

Do this instead

Start with a supervised assistant, not an autonomous theatre creature. Give it one narrow task, one bounded environment, logging, review, and a clear kill switch.

Field note

Fictionalised field note: A startup built an agent to manage outreach. It enriched contacts, drafted emails, scheduled follow-ups, and repeatedly contacted the same prospect under four names. The agent had not become sales. It has become a small automated social problem.

Takeaway

Autonomy without boundaries is not innovation. It is a liability with a cheerful interface.

5. “The model will figure it out.”

Claim

“The model will figure it out.”

What it pretends to mean

The system is intelligent enough to infer context, intent, edge cases, business rules, and the emotional weather of the finance department.

What it often means

No one wants to define the problem properly, so the model has been promoted to magical intern of ambiguity.

Proteus verdict

The model cannot read the meeting you failed to have.

Why it is dangerous

Modern models can generalise impressively, but they do not automatically know your business context, legal obligations, customer expectations, internal politics, data quirks, risk tolerance, or which spreadsheet column is named “final_FINAL_revised_USE_THIS”. If the task depends on hidden rules, tacit knowledge, or fragile exceptions, “the model will figure it out” becomes a polite way to outsource responsibility to probability.

The useful version

The useful version gives the model structured context, examples, constraints, rules, and evaluation. It treats the model as a system component, not a clairvoyant with a loading spinner.

Ask this before spending money

- What context does the model need to perform this task?
- What examples define a good output?
- Which edge cases matter?
- How will we test whether it figured out the right thing?

Do this instead

Build a small evaluation set before deployment: typical cases, edge cases, bad cases, ambiguous cases, and cases where the correct answer is “do not answer; escalate”.

Field note

Composite field note: A company used a model to classify customer complaints without defining escalation rules. The model politely categorised serious legal complaints as “general dissatisfaction”. It had figured something out. Unfortunately, it was the wrong civilisation.

Takeaway

If you refuse to define the task, do not be surprised when the model invents civilisation around the gap.

6. “We have data, so we have intelligence.”

Claim

“We have data, so we have intelligence.”

What it pretends to mean

Because the organisation has stored information, insight is now inevitable and possibly waiting near the dashboard.

What it often means

The data lake is a swamp with naming conventions, but everyone has agreed to call it strategic.

Proteus verdict

Data is not intelligence. It is memory with hygiene issues.

Why it is dangerous

AI systems depend on data quality, relevance, access, labelling, governance, context, and representativeness. Having data is not the same as having usable intelligence. Bad data can produce misleading outputs at scale. Fragmented data can create contradictory realities. Biased data can reproduce old decisions with new confidence. Unclear data ownership can turn a simple project into an archaeology expedition with legal consequences.

The useful version

The useful version begins by asking whether the data is available, lawful to use, relevant, clean enough, current enough, representative enough, and connected to the decision being made.

Ask this before spending money

- What data will the system use?
- Who owns it and who is allowed to access it?
- How clean, current, and representative is it?
- What important reality is missing from the data?

Do this instead

Run a data readiness inspection before the AI purchase: sources, permissions, quality, missing fields, known bias, update cadence, and who is responsible for corrections.

Field note

Fictionalised field note: A retailer built forecasting around customer data that excluded half its wholesale orders because those were tracked in an ancient spreadsheet called “misc”. The model was elegant. The stockouts were traditional.

Takeaway

Garbage in, garbage out. With AI, garbage may now arrive formatted.

7. “AI will replace judgement.”

Claim

“AI will replace judgement.”

What it pretends to mean

The system will make better decisions than humans, reducing bias, slowness, uncertainty, meetings, and the dangerous unpredictability of people having context.

What it often means

Management would like decisions without decision-makers, accountability without names, and confidence without awkward moral residue.

Proteus verdict

AI can assist judgement. Replacing judgement is how adults hide from consequences.

Why it is dangerous

Automation bias means people may over-trust system recommendations, especially under time pressure or uncertainty. The danger is not just that a model may be wrong. It is that humans may stop checking because the output feels official. Decision systems can also hide values inside metrics: what is optimised, what is ignored, who is harmed, and who gets to appeal. A model cannot be the final moral adult in the room.

The useful version

The useful version uses AI to inform, compare, flag, draft, calculate, retrieve, and surface patterns while preserving accountable human judgement for decisions with meaningful risk.

Ask this before spending money

- Which judgement is being assisted, and which remains human?
- What evidence must the human reviewer inspect?
- When must the decision be escalated?
- How can affected people challenge the outcome?

Do this instead

Write a judgement boundary: what AI may recommend, what it may not decide, who signs off, and what evidence the reviewer must examine before approving.

Field note

Composite field note: A hiring team used an AI ranking system and claimed humans were still in control. In practice, nobody challenged the top-ranked list because the numbers looked scientific and everyone was busy. The machine did not replace judgement. It sedated it.

Takeaway

Human-in-the-loop is not a decoration. The loop must contain a functioning human.

8. “The demo worked.”

Claim

“The demo worked.”

What it pretends to mean

The product has proven itself and is ready to enter the messy digestive system of the organisation.

What it often means

The system performed beautifully under conditions lovingly arranged by the people selling it.

Proteus verdict

A demo is a stage trick with Wi-Fi unless reality is invited to heckle.

Why it is dangerous

Demos are designed to reduce doubt. Real work is designed by nobody and contains exceptions, poor inputs, impatient users, weird integrations, missing data, edge cases, legacy systems, and human behaviour that looks like sabotage but is often just Tuesday. The danger is mistaking a controlled performance for operational evidence. A demo can show possibilities. It cannot prove durability in your environment.

The useful version

The useful version follows a demo with a pilot using your data, your users, your edge cases, your review process, your security constraints, and your success metric.

Ask this before spending money

→ Was the demo run on our data or vendor-selected examples?

- What failure cases were shown?
- How does it perform on messy inputs?
- What result would make us reject it?

Do this instead

Design a hostile pilot. Include bad inputs, ambiguous cases, edge cases, and a clear pass/fail threshold. If the vendor resists reality, thank them for the theatre and leave quietly.

Field note

Fictionalised field note: A procurement team approved a tool after a flawless demo. In deployment, the tool failed on regional formats, duplicate records, missing fields, and users who did not behave like demo actors. The demo remained perfect, safely preserved in memory like a tiny museum of optimism.

Takeaway

The demo is not the product. It is the product wearing stage makeup.

9. “AI saves time automatically.”

Claim

“AI saves time automatically.”

What it pretends to mean

The system will instantly reduce labour, accelerate work, and return human beings to meaningful creativity, strategic thinking, and possibly lunch.

What it often means

The task becomes faster while review, correction, integration, prompting, governance, exception handling, and anxiety quietly move elsewhere.

Proteus verdict

AI can save time. It can also move the work into a cupboard and call the cupboard innovation.

Why it is dangerous

Time savings are frequently claimed at the task level and ignored at the system level. A model may draft something quickly, but humans may spend time checking, editing, verifying, reformatting, integrating, explaining, defending, and repairing it. Hidden labour can erase the benefit. Worse, if AI increases output volume without increasing decision quality, the organisation may drown in more drafts, more options, more summaries, and more things nobody has time to inspect.

The useful version

The useful version measures end-to-end time: from input to accepted output, including review, correction, handoff, and rework. It also asks whether the saved time is actually removed, or merely filled with new noise.

Ask this before spending money

- Where exactly is time saved?
- What new review or correction time is created?
- Does output quality improve, hold, or degrade?
- What work will stop because of the saving?

Do this instead

Run a before-and-after time study on the whole workflow, not the impressive ten-second moment in the middle. Include review cost, rework, and downstream effects.

Field note

Composite field note: A marketing team adopted AI to produce campaign drafts faster. Output volume tripled. Review meetings quadrupled. Nobody saved time; they merely moved exhaustion from writing into sorting.

Takeaway

A faster draft is not a faster workflow if the humans now live in review hell.

10. “Everyone needs an AI strategy.”

Claim

“Everyone needs an AI strategy.”

What it pretends to mean

Every organisation must define a serious, forward-looking plan for AI or risk being crushed beneath the wheels of history and a vendor webinar.

What it often means

People want a strategy document before they have named the work, the customers, the risks, the data, or the reason the current spreadsheet is crying.

Proteus verdict

An AI strategy without a work strategy is a chandelier in a shed.

Why it is dangerous

AI strategy can be useful when it connects technology choices to real organisational goals. It becomes theatre when it floats above the work. Organisations then produce principles, committees, pilots, tool lists, and inspirational diagrams while the actual problems remain untouched. Strategy without prioritised use cases, governance, data readiness, capability building, evaluation, and budget discipline is not strategy. It is a weather forecast for procurement.

The useful version

The useful version begins with business reality: where work is slow, risky, repetitive, expensive, error-prone, knowledge-heavy, or blocked. Then it decides where AI might assist, what should not be automated, and how value and risk will be measured.

Ask this before spending money

- What organisational problem are we trying to solve?
- Which three use cases matter most and why?
- What capabilities, data, governance, and review do we need?
- Which AI uses are out of bounds?

Do this instead

Build an AI inspection map: problems, tasks, candidate use cases, risks, data readiness, human owners, pilot criteria, and stop rules. Do not start with a grand strategy. Start with work that can survive daylight.

Field note

Fictionalised field note: A firm commissioned an AI strategy before identifying any use cases. The final deck recommended innovation, enablement, transformation, and a centre of excellence. The invoices were highly real.

Takeaway

You do not need an AI strategy first. You need a reality strategy that knows when AI belongs in it.

What to Inspect Next

If this guide made an AI claim near you, begin sweating gently; further antidotes will appear in future Proteus materials published through Sienda Media. Proteus knows that there's more to be done. The fog has merely changed slides.

Use the questions. Inspect the workflow. Name the data. Find the reviewer. Price the review cost. Test the demo against reality. Make the adjective produce evidence before it receives money.

Approach Proteus when the pitch deck becomes too fragrant. Visit Sienda Media when you want the next field instrument. Bring gloves.

Sources, Further Reading, and Other Hype Detectors

This section is not here to turn satire into a seminar with weak coffee and twelve unread tabs. It is here because useful mockery should know what it is mocking.

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. Useful for thinking about AI as a risk-managed sociotechnical system rather than a magic sticker on a workflow.

OECD. (2019, updated 2024). OECD AI Principles. Relevant to trustworthy AI, human-centred values, accountability, transparency, and the radical idea that tools affecting people should not be operated like haunted vending machines.

European Union. (2024). Regulation (EU) 2024/1689, Artificial Intelligence Act. Included as a general regulatory context reference only. This guide is not legal advice; Proteus has no wig and refuses to bill in six-minute units.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? Proceedings of FAccT '21, 610-623. Important for grounding scepticism about language models, scale, data, bias, environmental cost, and the seductive mistake of treating fluent text as understanding.

Skitka, L. J., Mosier, K. L., & Burdick, M. (1999). Does automation bias decision-making? International Journal of Human-Computer Studies, 51(5), 991-1006. Useful for the warning that humans can over-trust automated recommendations, especially when the machine looks more confident than the room feels.

O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. New York: Crown. A practical reminder that data systems can scale harm when incentives, opacity, and weak accountability join hands in a dark corridor.

Mitchell, M. (2019). Artificial Intelligence: A Guide for Thinking Humans. New York: Farrar, Straus and Giroux. A clear, readable antidote to overestimating what AI systems understand, especially when they perform impressively on narrow tasks.

Marcus, G., & Davis, E. (2019). Rebooting AI: Building Artificial Intelligence We Can Trust. New York: Pantheon Books. Relevant to the limits of current AI systems and the need for more robust reasoning, common sense, and trustworthy design.

Kahneman, D. (2011). Thinking, Fast and Slow. London: Allen Lane. Useful background on confidence, judgement, bias, planning error, and the human habit of mistaking a neat answer for a correct one.

Frankfurt, H. G. (2005). On Bullshit. Princeton, NJ: Princeton University Press. A compact philosophical explosive for understanding speech that is not mainly committed to truth. Very useful near AI marketing.

Goodhart, C. A. E. (1975). Problems of Monetary Management: The U.K. Experience. In Papers in Monetary Economics, Volume 1. Reserve Bank of Australia. The source behind the metric warning often summarised as: when a measure becomes a target, it can stop being a good measure. Dashboards should be supervised accordingly.

Amershi, S., et al. (2019). Guidelines for Human-AI Interaction. Proceedings of CHI 2019. Useful for practical thinking about how humans and AI systems should communicate uncertainty, recover from errors, and keep users appropriately in control.