

The AI Theatre Detector

A practical field guide for small businesses, consultants, and managers
who need to spot expensive nonsense before the invoice arrives



Bob Mazzei

Sienda ltd

Copyright Notice

The AI Theatre Detector

A practical field guide for small businesses, consultants, and managers who need to spot expensive nonsense before the invoice arrives

© Sienda Ltd. All rights reserved. - 2026

Written by **Bob Mazzei**.

Published by **Sienda Media**, trademark of Sienda Ltd.

No part of this publication may be copied, reproduced, republished, uploaded, distributed, transmitted, resold, licensed, stored in a retrieval system, or used to create derivative commercial material without prior written permission from Sienda Ltd.

Brief quotations may be used for review, commentary, criticism, educational reference, or journalistic purposes, provided that appropriate attribution is given to the author and publisher.

This guide is provided for general educational and business-information purposes only. It does not constitute legal, financial, technical, medical, tax, investment, or professional advice. Readers should evaluate any AI tool, software proposal, consultant recommendation, or business process in the context of their own organisation, risks, obligations, and objectives.

All trademarks, product names, company names, and website names mentioned in this publication belong to their respective owners. Their mention does not imply endorsement, partnership, or affiliation unless expressly stated.

For permissions, enquiries, and further information, please contact Sienda Ltd through: sienda.co.uk or write to media@siendaweblines.com

About the Author

Bob Mazzei is an IT and Business Engineer, AI consultant, project manager, author, and founder of Sienda Ltd.

He has worked in information technology, business analysis, software projects, data analysis, and digital transformation for almost four decades. During his career, he has helped organisations understand, design, evaluate, and implement technological solutions across different sectors, always with a practical concern: technology should solve real problems, not create expensive theatre.

His work combines engineering discipline, business experience, philosophical inquiry, and a strong distrust of fashionable nonsense. He has helped companies implement AI when it made sense, and has advised them not to use it when it would add confusion, cost, disruption, or false sophistication without improving the business.

Bob's approach to AI is neither devotional nor reactionary. He sees artificial intelligence as a powerful technology, but not as magic, authority, or destiny. For him, the right question is not whether AI is "the future", but whether a specific tool solves a specific problem better, faster, safer, or more intelligently than the available alternatives.

Through Sienda Ltd, he develops and supports practical digital tools, AI-assisted business systems, secure information-exchange services, and market-intelligence solutions for small businesses, consultants, managers, and professionals.

The AI Theatre Detector was written to help business owners and decision-makers resist hype, ask better questions, protect their budgets, and distinguish useful AI from expensive performance art before the invoice arrives.

Acknowledgements

I am grateful to everyone who made this work possible: collaborators, friends, clients, and all the skilled people who help me, every day, do my work better.

My special thanks go to those who shared their thoughts on AI for the section *Notes from the Business Floor*. They are successful entrepreneurs who face complex problems every day, and who understand, from direct experience, the real difficulty of managing a business in the complicated landscape of our time.

Bob Mazzei

Table of Contents

Introduction	6
1 — The Great Fog Machine	11
2 — The Demo Trap	14
3 — The Consultant Incense Test	17
4 — The Agent Costume Department	20
5 — The Data Swamp	23
6 — The Workflow Autopsy	26
7 — The Invoice Detector	32
8 — The Five Questions That Ruin Bad AI Sales Calls	36
9 — When AI Is Actually Useful	39
10 — The Anti-Theatre Buying Checklist	42
11 — The Expertise Does Not Disappear	45
12 - Real World Cases.	52
Case #1	54
Case #2	55
Case #3	56
13 - Notes from the Business Floor	59
Appendix — The AI Theatre Detector Scorecard	66

Introduction

Welcome to the Theatre

One of the most exhausted slogans in the current circus is this: AI is here to stay.

It is usually delivered with the solemn expression of someone announcing a geological discovery, although very often by people who would struggle to explain what a byte is without injuring a metaphor. The sentence is treated as if it settles everything. AI is here to stay, therefore you must adopt it. AI is here to stay, therefore you must transform your business. AI is here to stay, therefore you must attend this webinar, buy this platform, hire this consultant, and allow a stranger with a slide deck to reorganise your operations around words like *agentic*, *autonomous*, and *disruptive*.

But the statement is almost meaningless.

Many things are here to stay. Stones, for example, have been here for billions of years and show no immediate sign of resigning from the planet. This does not mean every person must become a geologist. It does not even mean that builders, road workers, architects, or bridge engineers need to spend their professional lives worshipping rocks as a metaphysical revelation. They need to understand stones in relation to their actual work: their strength, their use, their limits, their behaviour under pressure, their cost, their risk, and their suitability for a specific purpose.

That is the sensible position with AI.

The question is not whether AI exists. It does. Very good. So do gravel, gravity, mould, taxation, and people who say “circle back” in meetings.

The real question is: what does this particular technology do for this particular business, in this particular process, under these particular conditions, with these particular risks and costs?

Anything else is theatre.

We should not fall into the frenzy of AI like irrational people dancing around a technological totem. But neither should we behave like the Luddites of the nineteenth century, smashing machines because history has once again committed the offence of moving.

Both attitudes are forms of fetishism.

The first worships the machine. The second fears it as a demon. One kneels before the new tool; the other throws stones at it. Both forget the same essential fact: technology is neither sacred nor cursed. It is made by human beings, used by human beings, misunderstood by human beings, exaggerated by human beings, and, quite often, sold by human beings with suspiciously glossy brochures.

Human beings are technical beings. This is not a recent development that began when someone opened a laptop in a co-working space and said “disruption” without shame. Our technical nature has a deep history, reaching back to earlier hominins who shaped tools, controlled fire, organised activities, transmitted skills, and slowly transformed survival into culture.

Even language itself can be understood as a kind of technology: a tool of thought, memory, coordination, persuasion, imagination, and social life. Every local idiom, every dialect, every professional vocabulary, every joke told in a village bar or boardroom, is part of this vast human machinery for organising experience.

So yes, AI is a technology.

But that is precisely the point.

It is a technology, not a miracle. It is a tool, not a deity. It may be powerful, useful, dangerous, overrated, transformative, clumsy, brilliant, expensive, absurd, or occasionally all of these before lunch. But it should be judged as every technology should be judged: by what it does, where it works, where it fails, what it costs, who controls it, what risks it introduces, and whether it solves a real problem better than the alternatives.

A hammer is not impressive because hammers are “the future”. It is impressive when there is a nail. A tractor is not valuable because industrial agriculture needed a mythology. It is valuable when it helps cultivate land more effectively. A spreadsheet is not sacred because it contains formulas. It is useful when it helps people calculate, compare, plan, and notice that someone has made a terrible mistake in column G.

AI deserves the same treatment.

No incense. No pitchforks. No mystical surrender. No theatrical panic.

Just the old human question that should haunt every tool before we pay for it:

What practical problem does this help us solve?

That question is the beginning of this book. Not “how do we join the AI revolution?” Not “how do we look modern?” Not “how do we avoid being left behind by a man on LinkedIn who has written the word *exponential* fourteen times before breakfast?”

The real question is simpler, colder, and much more useful:

What problem are we solving, how will we know it has been solved, and why is AI the right tool for this job?

This book exists because many small businesses, consultants, and managers are being pushed into AI decisions through fear, fashion, and fog. Some are told they are already late. Some are told their competitors are ahead. Some are told that every process must be automated, every employee must become “AI-native”, and every business must transform itself before the next invoice cycle or perish in a glorious bonfire of irrelevance.

Much of this is nonsense.

Some of it is expensive nonsense.

And a small business cannot afford expensive nonsense dressed as inevitability.

Large corporations can spend six months on a pilot, discover that it produced three dashboards, four internal committees, and no measurable value, then quietly bury the project under a rebrand. SMEs do not have that luxury. For a smaller organisation, a bad AI project is not a learning exercise. It is money removed from wages, marketing, stock, equipment, cash flow, and sleep.

This is why doubt matters.

Not ignorant doubt. Not lazy doubt. Not the reflexive suspicion of someone who wants every new tool thrown into the sea. The useful kind of doubt is disciplined doubt: the doubt of the scientist, the engineer, the philosopher, the serious practitioner. The person who checks assumptions, tests evidence, admits uncertainty, and changes position when reality refuses to co-operate.

This brings us to a small but revealing word: guru.

The word comes to us through Sanskrit, where it carried meanings such as weighty, serious, respectable, venerable. Behind it sits an older Indo-European root associated with heaviness. Originally, then, the guru was not a person with a ring light, a webinar funnel, and an urgent invitation to “transform your mindset before midnight”. The word suggested gravity. Weight. Authority. A person whose knowledge had substance.

Which makes its modern abuse rather magnificent.

Today, every discipline seems to produce its own small army of gurus: marketing gurus, productivity gurus, leadership gurus, crypto gurus, AI gurus. Many of them do not merely claim expertise. They radiate a faint expectation of reverence, as if the rest of us should approach their LinkedIn posts barefoot, carrying incense and a procurement budget.

This is where the trouble begins.

The attitude of the genuine scientist, philosopher, engineer, or serious practitioner is not theatrical certainty. It is disciplined doubt. The more complex the subject, the more careful the mind should become. Real knowledge does not strut into the room shouting that it has conquered reality. It

enters, examines the problem, checks the assumptions, tests the evidence, admits uncertainty, and revises its position when the facts refuse to behave.

That is not weakness. That is rationality.

We do not possess the inside of reality like a trophy in a glass cabinet. We approach it imperfectly, through models, measurements, arguments, experiments, mistakes, revisions, and better questions. Absolute certainty may be comforting, but comfort is not the same as truth. In fact, the worship of certainty is often the enemy of rational knowledge.

And rational knowledge is not an academic luxury. It is what allows us to fix real problems in business and in life.

When a person selling AI behaves like a guru, demanding belief before evidence, admiration before explanation, and payment before clarity, the warning light should come on.

Not because AI is useless.

Because reverence is a terrible procurement method.

1 — The Great Fog Machine

Making reality simplistic is always a mistake.

Unfortunately, it is also one of the most successful tricks in marketing.

Politics does it constantly. Technology does it with great enthusiasm. Business gurus do it as if complexity were a personal insult. A messy situation, full of constraints, history, people, incentives, costs, risks, habits, conflicts, exceptions, and boring operational details, is reduced to one shiny sentence.

AI will transform your business.

Automate everything.

Data is the new oil.

Work smarter, not harder.

Your competitors are already doing it.

These phrases work because they remove discomfort. They take complexity, put it in a blender, add confidence, remove evidence, and serve it in a glass called Strategy.

The result is often persuasive.

But persuasion is not proof.

A message can be successful because it is easy to repeat, not because it is true. A strategy can sell because it flatters the buyer, not because it solves the problem. A technique can spread because it sounds modern, urgent, and inevitable, not because it survives contact with invoices, employees, customers, suppliers, regulations, databases, and the immortal demon known as “the way we actually do things here”.

This distinction matters.

The fact that a marketing technique works does not mean the idea behind it deserves to be accepted. A slogan may win attention while still damaging judgement. A simplified promise may generate leads while hiding the very complexity that will later destroy the project.

And when the subject is not a campaign but the adoption of a business solution, the danger becomes much greater.

If a company buys into a simplistic idea of AI, automation, or digital transformation, the consequences are not limited to a mildly embarrassing LinkedIn post. Money is spent. Processes are disrupted. Employees lose time. Managers defend bad decisions because reversing them feels politically expensive. Vendors move on to the next miracle. The organisation is left with a tool nobody fully understands, a workflow nobody trusts, and a budget line that looks as if it was approved during a fever.

In marketing, simplification may attract attention.

In business operations, simplification can become a catastrophe wearing a nice font.

The point is not that every solution must be complicated. Good solutions are often simple. But they are not simplistic. They do not pretend that reality has fewer moving parts than it actually has. They simplify after understanding complexity, not before. They reduce friction without amputating the facts.

That is the difference between clarity and nonsense.

Clarity says: “Here is the real problem, here are the constraints, here is what we can improve, here is what we cannot promise, here is what success would look like.”

Nonsense says: “AI will change everything.”

One is useful.

The other is usually followed by an invoice.

The great fog machine works by making vague words feel concrete. It turns *transformation* into a product, *intelligence* into a sticker, *autonomy* into a sales promise, and *innovation* into a smell. It does not necessarily lie in a crude way. It does something more dangerous: it makes the buyer feel that doubt is backward, caution is cowardice, and asking for details is a failure to understand the future.

This is why the first skill of AI detection is linguistic hygiene.

Before asking whether a solution is good, ask whether the sentence describing it means anything.

When someone says a tool is “AI-powered”, ask: powered to do what?

When someone says it is “autonomous”, ask: autonomous within which limits, with what supervision, and who is responsible when it behaves like a caffeinated raccoon in a filing cabinet?

When someone says it “optimises workflows”, ask: which workflow, measured how, improved from what baseline, and for whom?

When someone says it is “seamless”, ask: seamless for the user, the manager, the vendor, or the person who will spend three weeks repairing the data afterwards?

Fog hates nouns. Fog loves abstraction. So the detector begins by forcing language back down to earth.

- Name the process.
- Name the cost.
- Name the risk.
- Name the owner.
- Name the output.
- Name the failure.
- Name the evidence.

If the proposal cannot survive that, it is not a strategy. It is atmosphere.

2 — The Demo Trap

A demo is not a system.

A demo is a performance under favourable lighting.

This does not mean demos are useless. A good demo can show potential, clarify an interface, help a buyer understand a workflow, and reveal whether a vendor has built something serious or merely trained a chatbot to smile at the budget. But the buyer must remember what a demo usually is: a carefully selected example, using prepared inputs, with a confident presenter, a controlled environment, and no angry customer asking why yesterday's delivery note has been classified as a poem.

The demo is where AI theatre shines.

The screen looks clean. The example is neat. The system appears to understand. The answer arrives in seconds. Everyone nods. Someone says "impressive". Someone else imagines the staff freed from drudgery, the reports writing themselves, the clients delighted, the invoices flowing, and the business finally operating like the elegant machine it has always failed to become because of people, reality, and Tuesdays or whichever day you are more inclined to blame.

Then the tool meets the actual business.

The real documents are badly scanned. The product names are inconsistent. The customer emails contain half a question, three attachments, and a complaint about something unrelated. The spreadsheet has hidden columns. The CRM has duplicates. The supplier changed its invoice format in March. The staff use different terminology. The process has exceptions that were never written down because everyone "just knows". The person who knows is on holiday.

This is where the demo dies.

Not necessarily because the technology is bad. Often, the demo dies because the demo was not asked to face the real environment.

A serious buyer should therefore never ask only: “Can you show me how it works?”

That is too easy.

The better question is:

Can you show me how it works with our material, our exceptions, our volume, our staff, our risks, and our least elegant Tuesday? Or whichever day you like to blame.

A staged miracle is not operational value. A chatbot that answers three perfect questions in a demo is not the same as a support assistant that survives a month of real customer messages. An invoice extractor that handles five clean PDFs is not the same as a system that processes hundreds of invoices from inconsistent suppliers and flags uncertainty honestly. A forecasting tool that looks clever on sample data is not the same as one that improves purchasing decisions without quietly inventing confidence.

The demo trap is not that demos are fake. The trap is believing that a demo proves the business case.

It does not.

A demo proves, at most, that something can look promising under selected conditions.

The buyer’s job is to move from performance to test.

A test uses real data. A test defines success before it begins. A test includes failure criteria. A test has a time limit, a budget limit, and a decision point. A test does not allow the vendor to move the goalposts every time the system disappoints reality.

For example, “Look how beautifully this AI summarises a meeting transcript” is a demo.

“Summarise ten of our actual sales calls, identify action points, flag uncertain sections, and let three managers compare the summaries against the recordings using agreed criteria” is a test.

“Look how this AI reads invoices” is a demo.

“Process one hundred of our real invoices, extract supplier name, VAT number, invoice date, total, line items, and purchase order reference, then show accuracy, exceptions, and human review requirements” is a test.

“Look how this AI assistant answers questions from the knowledge base” is a demo.

“Answer fifty real staff questions using our current documents, cite the source, admit when the answer is not available, and produce an audit trail” is a test.

A demo asks the buyer to admire.

A test asks the system to work.

The difference is where money should start becoming nervous.

3 — The Consultant Incense Test

A consultant can be useful.

That sentence must be said clearly before the satire puts on its boots. Good consultants exist. They bring perspective, structure, experience, technical knowledge, practical judgement, and the precious ability to tell a company that the problem is not the software but the process everyone has been politely avoiding since 2017.

The problem is not consultancy.

The problem is incense.

Consultant incense is the fog of vocabulary that fills a room when expertise is replaced by performance. It smells of strategy, transformation, disruption, and other words that have suffered too much presentation software exposure. It does not clarify the business problem. It elevates the atmosphere around it until everyone feels they are participating in something important, although no one can quite say what will be delivered by Friday.

AI has given consultant incense a new cathedral.

A weak consultant no longer needs to say merely that a company should “digitally transform”. Now they can talk about AI readiness, agentic workflows, autonomous decision layers, synthetic intelligence, proprietary orchestration, and exponential adoption curves. The words arrive dressed as expertise. Some may even be legitimate in the right context. But legitimacy in one context does not prevent abuse in another.

The Consultant Incense Test is simple: after the presentation, can the consultant describe the actual business problem in plain language?

Not the trend.

Not the opportunity.

Not the strategic landscape.

The problem.

If they cannot say, “Your sales team is losing leads because follow-up is inconsistent after the first contact,” or “Your invoice process fails because supplier names are stored in three different ways,” or “Your customer-service team needs classification and routing before it needs generation,” then the incense is winning.

A serious consultant should be able to move easily between technical language and operational language. They should understand the tool, but they should also understand the business situation in which the tool must live. They should not hide behind terms the buyer cannot evaluate. They should translate complexity without amputating it.

The warning signs are usually visible.

They speak about AI before they speak about the process.

They present a solution before diagnosing the problem.

They cannot identify who will maintain the system.

They do not ask what happens when the tool is wrong.

They avoid discussing data quality.

They treat integration as a minor detail.

They promise autonomy without explaining supervision.

They cannot define success in measurable terms.

They make doubt feel embarrassing.

That last sign matters.

A good consultant welcomes questions because questions reveal the shape of the work. A bad consultant treats questions as interruptions in the ceremony. If asking “what exactly will this do?” makes the room uncomfortable, the problem is not your ignorance. The problem is that someone has placed the invoice too close to the altar.

The antidote is not hostility. It is disciplined interrogation.

Ask them to explain the project without using fashionable vocabulary.

Ask them to identify the smallest useful pilot.

Ask them to show where AI is genuinely required and where ordinary software or automation would be enough.

Ask them what can go wrong.

Ask them what happens if they disappear.

Ask them what the company will own at the end.

Ask them how the result will be measured.

A consultant who can answer those questions may be useful.

A consultant who cannot still be entertaining, especially if the slides have gradients. But entertainment should be cheaper.

4 — The Agent Costume Department

Few words in modern AI marketing have acquired a costume faster than *agent*.

The word sounds active. Purposeful. Almost heroic. A chatbot sits there waiting for questions. An agent, by contrast, appears to stride into the business wearing a tiny operational cape, ready to book meetings, chase invoices, reorganise workflows, analyse customers, brief managers, update systems, and possibly make coffee if the roadmap survives Q3.

The problem is not that AI agents are imaginary. Agent-like systems can be useful. A well-designed assistant may take instructions, use tools, retrieve information, perform steps, draft responses, classify cases, update records, and ask for human approval when necessary. In the right setting, this can save time and reduce friction.

The problem is that the word *agent* is often used before the system deserves it.

Sometimes an “AI agent” is just a chatbot with access to a database.

Sometimes it is a form with better copy.

Sometimes it is automation with a language model bolted on the front like a theatrical moustache.

Sometimes it is a fragile sequence of prompts, integrations, and hopes, all described as autonomous until the first exception appears and the thing collapses into a puddle of apologies.

To understand agents, ignore the costume and inspect the behaviour.

What can the system actually do?

Can it take action, or only suggest action?

Can it use tools, or only produce text?

Can it access reliable data, or only improvise from the prompt?

Can it handle exceptions, or does it require perfect conditions?

Can it explain what it did?

Can a human review, approve, reverse, or stop it?

Can it fail safely?

That last question is critical.

An AI assistant that drafts an email badly is annoying. An AI system that sends the email without review may be dangerous. An AI assistant that suggests a supplier reorder is useful. An AI system that places the order automatically based on uncertain data may become expensive very quickly, especially if the warehouse begins receiving twenty-six pallets of the wrong thing because the agent felt confident.

Autonomy is not a decoration. It is a risk level.

The more action a system can take, the more supervision, logging, permissions, testing, and fallback it needs. A tool that drafts, suggests, or classifies may be acceptable with light review. A tool that changes records, approves payments, orders stock, contacts customers, or makes decisions affecting money, law, safety, or reputation requires serious control.

This is where the costume department becomes dangerous. Vendors love the glamour of autonomy because it sounds like value. Businesses should hear the word and immediately ask about brakes.

A good AI-assisted workflow might look like this:

A customer message arrives. The system classifies the topic, estimates urgency, extracts relevant details, suggests a response, and routes the case to the right person. A human reviews the suggestion, edits if necessary, and sends the reply. The system learns nothing mystical from this unless designed to do so; it simply helps organise work.

A dangerous version might look like this:

A customer message arrives. The system interprets it, generates a response, updates the CRM, promises a refund, changes the order status, and sends the reply without review. Everyone discovers the problem when the customer forwards the message to a lawyer, a journalist, or, worse, the managing director.

The difference is not whether both systems are “AI agents.”

The difference is control.

A useful agent is not one that pretends to replace judgement. It is one that supports judgement by handling defined tasks under defined limits.

Before buying an agent, ask:

What actions can it take, under what conditions, with whose approval, and how do we stop it when it is wrong?

If the vendor cannot answer, you are not looking at an agent.

You are looking at a costume.

5 — The Data Swamp

Most AI projects do not fail because the model is weak.

They fail because the business is feeding the model rubbish.

For small and medium-sized enterprises, this is one of the most painful discoveries. The sales conversation begins with intelligence, automation, efficiency, and transformation. The project begins with reality: spreadsheets, inboxes, paper receipts, WhatsApp messages, badly named folders, half-used CRMs, duplicated client records, inconsistent product names, and documents saved by people who believed “final-final-v3-real-this-one.xlsx” was an information architecture.

AI does not magically repair this mess.

It may magnify it.

If the input data is unclear, incomplete, contradictory, or badly structured, the result will be disappointing, misleading, or dangerous. A vendor may sell an SME an AI-powered system, only for the business to discover that months of cleaning, labelling, standardising, and organising data are required before anything useful can happen.

The model was not the first problem.

The cupboard was.

Data readiness is one of the least glamorous and most important parts of any AI project. It is also one of the easiest to hide during a sales pitch. Nobody wants to start a presentation by saying, “Before we discuss artificial intelligence, we need to investigate your folders, naming conventions, duplicates, missing fields, document formats, and the terrifying spreadsheet maintained by Annie.”

But that is often the work.

A system that classifies customer emails needs enough examples of real customer emails. A tool that extracts data from invoices needs invoices in formats it can read, with fields that can be checked. A

forecasting tool needs historical data that is consistent enough to support a forecast. An internal knowledge assistant needs documents that are accurate, current, accessible, and not contradicted by three forgotten policy PDFs from the reign of a previous manager.

The better question is not “Can AI use our data?”

The better question is:

Is our data fit to be used by anything more sophisticated than prayer?

This is not only a technical issue. It is an operational issue. Data reflects the way a business works. If responsibilities are unclear, data will be unclear. If processes vary by person, data will vary by person. If the company has never agreed what a lead, customer, order, product, complaint, or completed task actually means, AI will not repair the ambiguity. It will process it with magnificent speed.

The Data Swamp has several common creatures.

There is scattered data: information spread across too many places.

There is inconsistent data: the same customer, product, supplier, or process named in different ways.

There is missing data: fields left blank because everyone knows what they mean until suddenly nobody does.

There is outdated data: documents and records that appear official but no longer reflect reality.

There is unowned data: information nobody is responsible for maintaining.

There is decorative data: numbers collected because someone once liked the idea of reporting them, although nobody uses them to make decisions.

And then there is cursed data: the spreadsheet everyone depends on, nobody fully understands, and no one dares to open without first making a copy called “backup_real_final_do_not_touch”.

Before buying AI, inspect the swamp.

- Where is the data?
- Who owns it?
- Who updates it?
- How accurate is it?
- How complete is it?
- How often does it change?
- Can it be exported?
- Can it be checked?
- Does it describe the process as it really works, or as someone once hoped it worked during a planning meeting?

A serious vendor should ask these questions early. A serious buyer should welcome them, even if the answers are uncomfortable. The discomfort is useful. It prevents the business from paying for a model to discover what a fifteen-minute conversation with the operations team could have revealed.

AI can be useful with messy information. In fact, one of its strengths is handling unstructured text, varied documents, and imperfect human language. But useful does not mean miraculous. A model can assist with classification, extraction, summarisation, and search. It cannot turn organisational confusion into reliable intelligence merely by being expensive.

The Data Swamp must be mapped before it is automated.

Otherwise, the business does not get artificial intelligence.

It gets artificial confusion with a monthly fee.

6 — The Workflow Autopsy

Before automating anything, inspect the corpse.

That sentence sounds dramatic, but it is often the most humane advice a business can receive. Many workflows are not living processes. They are historical accidents with email addresses. They continue because someone inherited them, because a client once requested them, because a former manager designed them, because the software allowed them, or because everyone is too busy to ask whether the thing still needs to exist.

Then a vendor arrives and offers to automate them.

This is how companies make bad processes faster.

The intelligent question is not “Can we use AI?”

The intelligent question is:

Is this problem best solved by rules, by a model, by a human, by a better process, or by not doing the stupid thing in the first place?

To answer that, we need to distinguish ordinary software, automation, and AI.

Ordinary software follows instructions written by people. A calculator, a spreadsheet, an invoicing system, a booking platform, or an accounting programme may perform extremely useful and complex operations without being AI. It may calculate VAT, generate totals, apply formulas, create invoices, sort records, validate fields, and produce reports. This can save a business from chaos, bankruptcy, or a Friday afternoon nervous breakdown.

But usefulness is not the same as AI.

Automation is software arranged to perform a task automatically when certain conditions are met. It is the logic of “when this happens, do that”. When a customer fills in a form, send an email. When an invoice is marked as paid, update the accounting record. When stock falls below a

threshold, notify the purchasing manager. When a file is uploaded, rename it, store it, and send a message to the team.

That is automation.

It may be simple or complex. It may involve many steps. It may save hundreds of hours. But if the system is mainly following predefined rules, it is automation, not artificial intelligence in any serious modern business sense.

AI behaves differently. Instead of merely following explicit rules written in advance for every case, it uses a model trained on data to recognise patterns, classify inputs, predict outcomes, generate text, identify anomalies, extract meaning, or make recommendations. It does not usually say, “Rule 17 applies, therefore do X.” It says, in effect, “Based on what I have learned from many examples, this input most probably belongs to this category, means this, resembles that, or should produce this kind of output.”

Automation says:

When this happens, do that.

AI says:

Given this input, infer, classify, predict, generate, extract, rank, recognise, or recommend.

That difference matters.

A customer fills in a contact form. The system sends a confirmation email, adds their details to a spreadsheet, notifies the sales team, and creates a task in the CRM. That is automation. Useful? Yes. AI? No. Magic? Only if your previous office system involved Post-it notes, panic, and someone called Sandra remembering everything.

Now imagine a system that reads incoming customer messages, detects the subject, estimates urgency, classifies the request, extracts names, dates, order numbers, and complaints, suggests a reply, and flags cases that may require human attention. That may involve AI. Not because it sends

an email. Sending an email is not artificial intelligence. The AI part is the classification, extraction, interpretation, summarisation, or generation performed on messy human language.

The same distinction appears in invoice processing.

If a system receives an invoice from a known supplier, saves the file in the correct folder, sends it to accounting, and marks it as received, that is automation.

If a system reads the invoice, identifies supplier data, extracts line items, detects inconsistencies, compares it with purchase orders, flags unusual amounts, and suggests whether it should be approved or reviewed, then we may be looking at AI-assisted processing.

Again, the difference is not glamour. It is function.

A neural network is one way of building this kind of AI model. Very simply, it is a computational system inspired, very loosely, by the structure of biological nervous systems. It contains layers of artificial “neurons” that process signals. During training, the system adjusts internal numerical values, often called weights, so that it becomes better at producing the desired output from given inputs.

For example, a neural network may be trained on thousands or millions of images so that it learns patterns associated with cats, invoices, tumours, road signs, defective products, or handwritten numbers. It is not given a neat little rule such as “a cat is an animal with triangular ears, whiskers, moral superiority, and the legal right to sit on your keyboard”. Instead, it gradually learns statistical patterns from examples.

This is very different from a standard automation process.

An automation process says: “If the invoice amount is greater than £5,000, send it to the finance director.”

A neural-network-based system might say: “This invoice looks unusual compared with previous invoices because the supplier name, amount, timing, formatting, and purchase order relationship resemble cases previously flagged for review.”

The first is a fixed rule.

The second is pattern recognition.

Both may be useful. Both may be combined. But they are not the same thing.

This distinction is not academic. It is commercial self-defence.

If a supplier sells automation as automation, fine. The question is whether it works, saves time, reduces friction, and costs less than the problem it solves.

If a supplier sells automation as AI, the warning light should begin to flash. Not because automation is bad, but because the language is already being inflated. And inflated language is often followed by inflated invoices.

A business owner, consultant, or manager should ask a brutally simple question:

Which part of this solution is actually AI?

Not “Is this powered by AI?”

Not “Does this use AI somewhere?”

Not “Is there an AI layer?” — that horrible phrase, as if intelligence had been spread on the system like butter on toast.

Ask instead:

Show me the exact task where the system is not merely following a fixed rule, but making a prediction, classification, extraction, recommendation, generation, or interpretation based on data.

If the answer is vague, the product may still be useful, but the AI claim is probably theatre.

This also explains why the word intelligent causes so much trouble.

There is no single definition of intelligence accepted by everyone across philosophy, psychology, neuroscience, cognitive science, artificial intelligence, and everyday language. The word shifts depending on context. Human intelligence includes reasoning, adaptation, learning, abstraction, memory, language, planning, creativity, social judgement, bodily skill, emotional understanding, and much more. Some animals show forms of intelligence. Human institutions sometimes show none at all, despite having committees.

So when we call a machine *intelligent*, we must be careful.

If a calculator solves a complex equation, is it intelligent?

In one loose everyday sense, we might be tempted to say yes, because it performs a task that requires intelligence when done by a human. But in a stricter sense, the calculator is not intelligent. It is executing formal operations designed by human beings. It does not understand the equation. It does not know what a number is. It does not wonder whether mathematics reveals the structure of reality. It does not finish the calculation and then stare through the window, troubled by Pythagoras.

It calculates. And that is enough.

The same is true of many excellent business tools.

An invoicing system that calculates VAT correctly is useful.

A stock-control system that warns when inventory is low is useful.

A spreadsheet that models different pricing scenarios is useful.

A scheduling system that prevents double bookings is useful.

A database that retrieves the right client record in two seconds is useful.

None of these things need to be called AI to deserve respect.

A solution does not become better because it borrows the vocabulary of cognition. It becomes better when it works. The real intelligence may not be inside the software. It may be in the decision not to overbuy it.

7 — The Invoice Detector

Building an AI system like ChatGPT is not the same as adding a chatbot to a website, connecting an API, or asking a programmer to “put some AI in the app”, as if intelligence were a decorative plugin available between the cookie banner and the contact form.

A system of that kind requires several enormous layers working together: massive quantities of training data; teams of machine-learning researchers, software engineers, data engineers, infrastructure specialists, security experts, product designers, and safety reviewers; vast computing power using specialised hardware such as GPUs (Graphics Processing Units, specialised electronic chips designed to process vast amounts of data simultaneously at high speed) or other accelerators; data-cleaning pipelines; model-training infrastructure; evaluation systems; fine-tuning processes; monitoring tools; abuse prevention; privacy controls; legal review; constant maintenance; and the ability to serve millions of requests reliably, quickly, and safely.

This is not a normal business IT project.

It is industrial-scale research and engineering.

A company may reasonably build an AI-powered application. It may use existing models through APIs, customise workflows, connect internal documents, create specialised assistants, automate support, summarise reports, classify emails, extract data from invoices, or build practical tools around AI capabilities. That can be valuable.

But building the underlying foundation model — the thing that behaves like ChatGPT itself — is another matter entirely. It means creating, training, testing, deploying, and maintaining a general-purpose model with billions of parameters and a heavy operational infrastructure behind it. For most companies, this would be economically irrational, technically excessive, and strategically unnecessary. Or pure madness, if you prefer.

The difference is like the difference between using electricity and building a national power grid.

A small business does not need to construct a power station to turn on the lights. A consultancy does not need to manufacture its own laptop to write a report. A restaurant does not need to build the entire internet to accept online bookings. In the same way, most businesses do not need to build their own AI infrastructure. They need to understand where AI can help, choose the right existing tools, protect their data, define useful workflows, and measure whether the result actually improves anything.

This is where many AI theatre productions become absurd.

Someone says: “We should build our own AI.”

Fine. What exactly does that mean?

Do they mean building a full foundation model from scratch? Almost certainly madness, unless the company owns a mountain of money, a research division, serious data assets, specialised infrastructure, and a strong reason to compete at the frontier of AI development.

Do they mean building a business tool that uses an existing AI model? That may be perfectly sensible.

Do they mean connecting a chatbot to company documents? That is not “building an AI”. It is building an application that uses AI.

Do they mean automating part of a workflow with some AI-assisted classification or text generation? Again, possible, useful, and far more modest than the slogan suggests.

The practical point is simple: most companies should not try to build “their own ChatGPT”. They should build or buy carefully designed tools that use AI where AI genuinely adds value. The intelligence of the business decision lies not in pretending to become an AI laboratory, but in knowing which layer of the technology stack belongs to you and which layer should be rented, integrated, governed, and kept under adult supervision.

A company does not become more advanced by rebuilding what specialists have already built at astronomical cost.

It becomes more advanced by solving its own problems with proportion, clarity, and discipline.

For SMEs, this matters because AI costs are often less obvious than ordinary software costs.

AI is rarely a simple “buy once and relax” purchase. Many AI services carry ongoing costs: API usage, cloud infrastructure, storage, monitoring, maintenance, security, updates, support, retraining, and technical debt. What begins as a friendly £5,000 project can quietly become a monthly subscription monster wearing a tiny hat.

This matters because SMEs operate under real constraints. A large company may absorb unpredictable costs. A small company may not. If a tool depends on token usage, cloud compute, premium tiers, external APIs, or vendor-controlled pricing, the business must understand what happens when usage grows. Otherwise, a project that looked profitable in the sales meeting may become absurd in month three.

Skill dependency is another cost, although it is not always written clearly on the proposal.

SMEs usually cannot afford a proper data science team. They may rely on one consultant, one vendor, or one enthusiastic employee who has watched six videos and now says “fine-tuning” with priestly confidence. If that person leaves, disappears, changes prices, or turns out to have confused automation with AI, the business is exposed.

This creates a dangerous imbalance. The vendor speaks in technical language; the owner cannot easily judge the work; the system becomes a black box; and soon the company is not buying a solution but renting helplessness.

Then comes lock-in.

A vendor who refuses to provide export options, documentation, access to your data, or a reasonable exit path is not selling technology. They are building a small digital prison and asking you to admire the curtains.

At minimum, an SME should know whether it can export its data in ordinary formats such as CSV, JSON, PDF, or plain document files. It should know whether workflows can be documented. It should know whether another provider could maintain the system. It should know what happens when the contract ends.

The most dangerous AI vendor is not always the most expensive one.

It is the one that makes leaving impossible.

So before approving an AI invoice, ask:

- What are the upfront costs?
- What are the monthly costs?
- What costs grow with usage?
- What happens if usage doubles?
- Who maintains the system?
- Who pays for fixes?
- Who owns the data?
- Can we export everything?
- Can we leave without chaos?
- What will we still need to pay for six months from now?

A proposal that cannot answer these questions is not ready for approval.

It may be ready for theatre.

But theatre tickets should not be filed under operational investment.

8 — The Five Questions That Ruin Bad AI Sales Calls

Bad AI sales calls survive on momentum.

The vendor presents. The slides advance. The language becomes larger. The future arrives in gradients. The buyer begins to feel that asking a simple question might reveal ignorance. Everyone nods, because nodding is socially safer than interrupting a man who has just used the phrase “proprietary intelligence layer.”

This is how bad projects begin.

The antidote is not aggression. It is clarity.

The following five questions ruin bad AI sales calls because they force the conversation away from atmosphere and towards reality.

1. What exact business problem does this solve?

Not “what trend does it belong to?” Not “what opportunity does it represent?” Not “how exciting is this technology?”

What problem?

“We need AI” is not a problem. “We spend eight hours per week copying invoice data from PDFs into accounting software” is a problem. “We lose leads because follow-up is inconsistent after the first enquiry” is a problem. “Customer emails are not sorted by urgency, so important cases are buried” is a problem.

If the vendor cannot name the problem in plain business language, the solution is premature.

2. Which part of this is actually AI?

This question is wonderfully impolite to inflated language.

It forces the vendor to distinguish ordinary software, automation, and AI. Is the system following fixed rules? Is it moving data from one place to another? Is it sending scheduled notifications? Is it applying a template? Or is it genuinely classifying, predicting, extracting, ranking, generating, recognising, or interpreting based on data?

The answer matters because automation may be exactly what the business needs. But if automation is being sold as AI, the price and risk may already be swollen.

3. What happens when it is wrong?

Every serious AI discussion must include failure.

Not if it is wrong. When.

Does the system flag uncertainty? Does a human review important outputs? Is there an audit trail? Can a decision be reversed? Can a user see the source? What happens if the model misclassifies a customer complaint, extracts the wrong invoice amount, invents a confident answer, or sends something inappropriate?

A vendor who has no answer to failure has no serious system.

They have a demo with optimism attached.

4. Who cleans the data, owns the workflow, and maintains the system?

This question attacks the hidden labour.

AI projects do not live in the abstract. They live in data, documents, integrations, permissions, updates, staff habits, business rules, exceptions, and maintenance. Someone must clean the data. Someone must decide which documents are current. Someone must update prompts, rules, workflows, integrations, and permissions. Someone must handle errors.

If the vendor says, “You do”, then they must explain how, with what tools, what training, and how much time.

If the vendor says, “We do”, then they must explain the cost, contract, limits, and handover.

If nobody knows, the project already has a ghost invoice waiting inside it.

5. How do we leave?

This is the question many buyers forget to ask because the sales call is still in its honeymoon period.

Can the company export its data?

Can it export prompts, configurations, workflows, documents, logs, and outputs?

Are there exit fees?

Can another provider maintain the system?

What happens if the vendor disappears, changes pricing, removes a feature, or becomes unbearable?

A serious vendor will not be offended by exit questions. They know that trust increases when the buyer is not trapped.

A weak vendor may become vague.

Vagueness here is not a detail.

It is a warning.

These five questions do not guarantee success. Nothing does. But they remove much of the fog. They make bad vendors uncomfortable and good vendors more useful. They also remind the buyer of the central discipline of this book:

Do not buy AI.

Buy a measurable improvement to a real business problem, using AI only where AI genuinely earns its keep.

9 — When AI Is Actually Useful

This book spends a lot of time mocking AI theatre because AI theatre deserves to be mocked. It wears expensive shoes, speaks in fog, and keeps placing invoices where evidence should be.

But the point is not that AI is useless.

The point is that usefulness must be earned.

AI can be genuinely valuable when the task involves messy information, language, images, patterns, uncertainty, classification, extraction, prediction, generation, summarisation, or decision support. It is especially useful when the problem is narrow enough to test, valuable enough to matter, and safe enough to supervise.

The safest early opportunities for SMEs are usually not glamorous. This is good news. Glamour is often where budgets go to die.

Useful AI often begins with ordinary irritation.

Customer emails arrive in different styles, languages, levels of anger, and degrees of punctuation collapse. AI can help classify them by topic, urgency, sentiment, product, client type, or required action.

Invoices arrive as PDFs, scans, attachments, and small acts of supplier creativity. AI can help extract supplier names, dates, totals, VAT numbers, purchase order references, and line items, while flagging uncertainty for human review.

Sales calls and meetings produce long recordings nobody wants to revisit unless legally required. AI can transcribe, summarise, identify action points, extract objections, and prepare follow-up drafts.

Internal documents multiply until finding the right policy becomes a spiritual trial. AI-assisted search can help staff ask questions in natural language and receive answers grounded in approved documents.

Support teams repeat similar explanations. AI can draft replies, suggest knowledge-base articles, and help new staff understand previous cases, provided humans remain responsible for sensitive responses.

Managers need first drafts of reports, proposals, emails, summaries, and briefing notes. AI can accelerate drafting, restructuring, comparison, and editing. It does not remove judgement. It gives judgement something to work on before lunch.

Operations teams need anomaly detection. AI can help spot unusual transactions, unexpected patterns, sudden changes, missing fields, repeated errors, or cases that deserve review.

Marketing teams need variations, outlines, summaries, translations, and adaptation of content. AI can help, provided someone with taste, strategy, and a functioning relationship with reality remains in charge.

In all these cases, the value does not come from the word AI.

It comes from reducing time, error, friction, delay, or missed opportunity.

A useful AI project usually has several traits.

- The problem is specific.
- The input is available.
- The output can be checked.
- The risk is understood.
- The workflow is clear.
- The human role is defined.
- The cost is visible.
- The benefit can be measured.
- The system can fail safely.

This is why AI is often most useful as assistance rather than replacement. It drafts, suggests, classifies, searches, extracts, compares, flags, and summarises. The human still decides, approves,

corrects, owns, and notices when the machine has produced something that sounds confident but has the structural integrity of wet cardboard.

The best early AI use cases for SMEs are often narrow, boring, and profitable.

Excellent.

Boring and profitable is a noble category. Many businesses would benefit from more boring and profitable, and fewer visionary initiatives requiring a launch video.

The right question is not “Where can we use AI?”

That question invites theatre.

The right question is:

Where do we repeatedly spend time, money, or attention on work that involves messy information and could be assisted, checked, and improved?

That is where AI may deserve a seat at the table.

Not the throne.

A seat.

Preferably away from procurement until it has passed the test.

10 — The Anti-Theatre Buying Checklist

For SMEs, the intelligent path is not to chase the largest promise. It is to find the smallest valuable problem, test it with real data, measure the result, control the risk, and scale only when the evidence deserves it.

The first rule is simple: never begin with AI. Begin with a painful, specific, repeated business problem.

Not “we need AI”.

Instead:

“We spend six hours per week copying invoice data.”

“We lose leads because follow-up is inconsistent.”

“We cannot compare supplier prices quickly.”

“We receive customer emails that need sorting by urgency.”

“We waste time rewriting similar proposals.”

“Our stock information is always late.”

Now the conversation is sane. Now AI, automation, ordinary software, a spreadsheet, a better form, or a human process change can compete honestly. The winner should be the method that solves the problem most safely, cheaply, and reliably.

The second rule is to demand a pilot with a hard stop.

An SME should be very careful with vendors who want a large upfront payment for a six-month custom AI build wrapped in fog. A serious vendor should be able to define a small test using real business material, anonymised if necessary, with written success criteria.

For example: “Process one hundred invoices and extract supplier name, VAT number, invoice date, total amount, and purchase order reference with at least 95% accuracy, while flagging uncertain cases for human review.”

That is a test.

“Transform finance operations with AI-driven intelligence” is a hymn sung before the payment.

The pilot should have a time limit, a budget limit, and a decision point. If the pilot fails, the business stops. No emotional blackmail. No “we are almost there”. No “the next phase will unlock the value”. The purpose of a pilot is not to become a cage with nicer paint.

The third rule is never to buy “AI” as such. Buy a measurable business outcome.

Do not ask: “Does it use AI?”

Ask:

Will this reduce manual work by ten hours per month?

Will this cut response time from two days to two hours?

Will this reduce invoice errors?

Will this help us quote faster?

Will this help us find missed opportunities?

Will this make reporting more accurate?

Will this save more money than it costs?

If the answer cannot be expressed in practical terms, the vendor is probably selling atmosphere.

The fourth rule is to keep a manual override for critical processes.

Never place payroll, legal decisions, inventory ordering, credit approval, safety procedures, compliance responses, or customer-critical decisions entirely in the hands of AI without human supervision. Full autonomy sounds impressive until the system confidently does something mad and everyone begins searching the contract for the word “liability”.

AI can assist. It can classify, draft, detect, summarise, compare, extract, and recommend. But in critical processes, there must be a human checkpoint, an exception route, and a way to stop the machine without setting fire to the business.

The fifth rule is to beware of lock-in.

A vendor who refuses to provide export options, documentation, access to your data, or a reasonable exit path is not selling technology. They are building a small digital prison and asking you to admire the curtains.

The sixth rule is to verify simply, quickly, and repeatedly.

Never pay for an AI outcome you cannot verify in a practical way. That does not mean every AI result must be perfect. It means the business must have a way to check whether the tool is doing anything valuable. If the claimed benefit cannot be tested with a small sample, a spreadsheet, a manual comparison, or a short operational experiment, the project may be too vague, too risky, or too theatrical for an SME to buy safely.

The checklist is not designed to make the buyer cynical.

It is designed to make the buyer dangerous to nonsense.

11 — The Expertise Does Not Disappear

One of the most persistent misconceptions about AI is that it will replace expertise.

This is convenient, dramatic, and commercially attractive. It allows vendors to imply that a company can buy a tool and immediately reduce its dependence on trained people. It allows managers under pressure to imagine a future where awkward human costs, professional judgement, experience, training, and responsibility can all be compressed into a subscription plan. It allows the excitable prophet class to announce that nobody needs to learn anything any more, because the machine will do it all.

This is not only wrong.

It is dangerous.

AI does not abolish expertise. In many contexts, it increases the need for it.

Consider medicine.

I am not a physician. If an AI-powered medical machine analyses images, scans, signals, probabilities, or patient data, the existence of AI does not suddenly make me qualified to interpret the result. The machine may assist diagnosis. It may detect patterns. It may highlight anomalies. It may support clinical decision-making. But the interpretation belongs to trained medical professionals who understand anatomy, pathology, symptoms, risk, context, patient history, uncertainty, false positives, false negatives, and the terrifying fact that behind every data point there is a human being who would prefer not to become an example in a negligence seminar.

The AI system may be powerful.

The responsibility remains human.

The same principle applies in business.

An AI tool can assist with customer support, legal drafting, financial analysis, marketing copy, invoice processing, stock forecasting, recruitment screening, sales communication, and internal knowledge search. But assistance is not a replacement. A tool that produces an answer does not automatically make the person reading the answer competent to judge it.

This is where many companies become reckless.

They imagine that because AI can produce language, summaries, calculations, classifications, designs, suggestions, and reports, the expertise required to evaluate those outputs has somehow vanished. It has not vanished. It has moved. Sometimes it has moved from production to supervision. Sometimes from execution to verification. Sometimes from writing the first draft to knowing whether the first draft is dangerous nonsense with excellent punctuation.

If you cannot write, AI will not make you win the Nobel prize for literature.

It may produce sentences. Some may even be elegant. But writing is not only the generation of sentences. Writing involves judgement, taste, structure, rhythm, purpose, audience, revision, proportion, voice, truthfulness, and the delicate ability to remove the sentence you loved because it is showing off like a drunk peacock.

AI can help a good writer move faster. It can help a weak writer produce something less terrible. But it does not magically install literary judgement in the user.

If you cannot paint, AI will not make you Michelangelo.

It may produce an image. It may even produce a beautiful image. But visual creation is not only the production of pixels. It involves composition, anatomy, light, colour, symbolism, style, intention, constraint, and the ability to see when something is wrong. Without artistic judgement, the user may admire a hand with seven fingers because the background is dramatic.

Technique alone does not become art by magic.

This does not mean that AI has no place in new artistic horizons. It clearly does. There are already artists, designers, musicians, filmmakers, writers, and visual experimenters using AI in intelligent, original, and fascinating ways. In those cases, AI becomes part of the instrumentarium of creation: a new tool, a new surface, a new collaborator, a new provocation, a new way of exploring form.

That is legitimate.

But it is very different from saying that AI gives the user artistic mastery.

No AI will ever hand you a Renaissance sculptor's skill merely because you have written a prompt. If an AI-powered machine were to produce an excellent sculpture, even a technically astonishing one, it would not have the same value as Michelangelo's *Pietà*. The reason is not simply that one is old and the other is new, or that one is made of fine marble and the other perhaps of a less noble material. The reason is that the human skill, vision, discipline, suffering, imagination, hand, eye, and spiritual tension involved are not remotely comparable.

The *Pietà* is not merely an object with a beautiful form. It is the visible trace of a human mastery that passed through body, mind, material, tradition, faith, anatomy, technique, and something far beyond technique. It carries the weight of a gesture that was not mechanically produced but artistically incarnated.

The same applies to every form of creative expression: literature, painting, sculpture, music, architecture, theatre, cinema, dance, and all the strange hybrid forms still being born.

Technique can be learned. It can be practised, refined, disciplined, and perfected. I can learn the rules of composition. I can study anatomy. I can practise drawing. I can learn harmony. I can train my hand, my ear, my eye, my rhythm. I can learn to play the violin.

But to be Paganini is something else.

Paganini was not merely a person who operated a violin with advanced technical competence. His art began where technique ceased to be visible as technique. The gesture became music. The

discipline disappeared into expression. The hand no longer seemed to execute; it seemed to reveal. The technical act was surpassed in the very act of being performed.

That is one of the mysteries of art.

The artist must learn technique, but true art is not the exhibition of technique. It is the moment when technique is absorbed, forgotten, transcended, and delivered into beauty. The gesture becomes so deeply mastered that it no longer appears as effort. It becomes form. It becomes necessity. It becomes presence.

AI may imitate styles, generate variations, compose plausible structures, produce images, assist musicians, suggest phrases, and open creative pathways. It may even help artists discover things they would not have found alone.

But it does not suffer the discipline. It does not inhabit the gesture. It does not know the resistance of marble, the terror of the blank page, the years of failure behind a single line, the physical intimacy of the instrument, the silence before the note, the judgement that rejects ninety-nine possibilities before choosing one.

This is why confusing AI output with artistic mastery is another form of theatre.

The tool may be powerful.

The result may be interesting.

The experiment may be worthwhile.

But the possession of a tool is not the possession of art. And the generation of an object is not the same as the human conquest of form.

With the right term of comparison, the same applies to management, law, accounting, engineering, sales, medicine, education, and strategy. AI can generate, suggest, and accelerate. It cannot remove the need to know what good looks like.

This is why the slogan “*You do not need to be an expert any more*” is one of the most poisonous forms of AI theatre.

It sounds democratic. It sounds empowering. It sounds like liberation from credentialed gatekeepers. Sometimes, in a limited sense, AI really does lower barriers. It can help people draft, learn, explore, translate, summarise, prototype, and understand topics that once required more time or support. That is valuable.

But lowering the first barrier is not the same as eliminating the mountain.

A beginner can now produce more. That does not mean the beginner can judge more. And in professional contexts, judgement is often the thing being paid for.

This matters especially inside businesses.

It is not advisable to implement AI in a business process without adequately training the staff who will use it. Dropping AI into an organisation without training is not modernisation. It is releasing a clever animal into the office and hoping it reads the employee handbook.

Staff need to understand what the tool can do, what it cannot do, where it may fail, what data it can use, what data it must not use, when to trust it, when to check it, when to escalate, and how to document decisions. They also need to understand the boundaries of responsibility. If the AI drafts a customer response, who approves it? If it flags an invoice as suspicious, who investigates? If it summarises a legal document, who verifies the meaning? If it produces a report, who owns the conclusion?

A company that trains nobody is not implementing AI.

It is outsourcing confusion to a user interface.

Even using a general tool like ChatGPT at a professional level requires application and study. Many people believe that because the interface is simple, the skill must be simple. This is another mistake.

A piano also has visible keys. This does not mean the user is now a pianist because they have located middle C.

Professional AI use requires learning how to frame a task, provide context, define constraints, evaluate output, check sources, identify hallucinations, compare alternatives, protect confidential information, refine instructions, and recognise when the tool is not suitable. It requires practice. It requires domain knowledge. It requires scepticism. It requires humility to understand that a fluent answer is not necessarily a correct answer.

AI is not a magic solution poured over incompetence.

If the process is confused, AI may accelerate the confusion.

If the data is poor, AI may process rubbish elegantly.

If the user lacks judgement, AI may provide confidence without competence.

If the organisation has no training, AI may distribute risk faster than anyone can manage it.

This does not make AI useless. It makes implementation serious.

The companies that benefit from AI are not necessarily those that shout about it the loudest. They are the ones that connect it to real problems, train their people, define responsibility, preserve human judgement, and understand that tools amplify the capabilities of the system into which they are placed.

A good doctor with a good diagnostic tool may become more effective.

A good accountant with AI-assisted document review may save time.

A good customer-service team with AI classification may respond faster.

A good manager with AI-supported reporting may see patterns sooner.

A good writer with AI assistance may draft, test, and revise more efficiently.

But the word **good** matters.

AI amplifies. It does not sanctify.

It can amplify skill, but it can also amplify carelessness. It can accelerate insight, but it can also accelerate malarkey. It can support expertise, but it cannot replace the responsibility of knowing what one is doing.

So the final misconception to reject is this: AI will let businesses avoid expertise.

No!

AI will reward businesses that respect expertise enough to use the tool properly.

The future does not belong to people who blindly worship AI, nor to people who reject it like frightened villagers watching a steam engine. It belongs to people who know enough to ask better questions, train their teams, keep humans responsible, and treat AI as a powerful instrument rather than a substitute for thought.

The machine may help.

The judgement is still ours.

Before the invoice arrives, the question is not whether AI sounds impressive. The question is whether the proposal survives examination. The scorecard below is designed for exactly that moment: the meeting has ended, the slides are still glowing faintly in memory, and someone is about to approve spending money. Use it before the fog becomes a contract.

12 - Real World Cases.

My first encounter with neural networks dates back to the early 1990s. I was introduced to them by a remarkably talented and knowledgeable software engineer, someone who understood that this was not simply another programming technique, but a different way of thinking about computation itself.

At the time, of course, there were no convenient frameworks such as TensorFlow or PyTorch. There were no friendly libraries capable of hiding the mathematical and computational complexity behind a few lines of elegant code. We worked much closer to the machine, mainly with C and Fortran. Fortran — launched by John Backus in December 1953 — in particular, had a long and distinguished history in scientific computing. It was built for heavy numerical work, matrix operations, simulations, and all those demanding calculations that made it a natural environment for early experiments in neural networks, climate modelling, financial forecasting, engineering analysis, and similar fields.

I used Fortran myself in that period, and I remember enjoying it immensely. It was only a pilot project, mainly involving simulations, but for me it was intellectually exciting. It had very little to do with the kind of software work I had done in the previous years, yet precisely for that reason it represented a valuable leap in my technical and cultural formation. It opened a door onto a different landscape: one where software was not merely instructed step by step, but could be trained, adjusted, and made to recognise patterns through examples.

In a career in IT that now spans nearly thirty-eight years, I have witnessed many technological revolutions. I saw the spread of personal computers, the rise of networks, the explosion of the internet, the transformation brought by mobile devices, cloud computing, and countless waves of software innovation. Yet I do not think even the mass diffusion of the first PCs felt quite as impressive as the arrival of mass AI today.

And this is the important point: AI did not suddenly appear in the 2020s, as some fashionable narratives seem to suggest. Its roots go back to the 1950s, and many of the ideas behind today's

systems have been studied, tested, forgotten, revived, and refined over decades. What is new is not the existence of artificial intelligence itself, but its scale, accessibility, computing power, and everyday presence. The technology has moved from laboratories, specialist environments, and experimental projects into the hands of the general public.

That is what makes the current moment extraordinary. Not that AI was invented yesterday, but that something with such a long and complex history has finally become a mass phenomenon.

So, once again, we should be cautious about all this hype. Much of it makes little sense and often feels less like serious technological discussion than a kind of fetishism.

Therefore, although the term *artificial intelligence* was coined in 1956, during the Dartmouth Summer Research Project on Artificial Intelligence, in the years when I was working with neural networks people generally spoke of “neural networks”. Almost no one in the field routinely used the expression *artificial intelligence*, and it was certainly not common among non-specialists.

There were no clients asking for “accounting software powered by neural networks”. There were no marketing departments dressing up ordinary software with fashionable AI language. Neural networks were not being sold as a universal miracle, a magic wand, or the latest shiny novelty to be attached to every product brochure.

That is why today’s situation is so different. The decisive change is not simply the technology itself, but its mass diffusion, its commercial packaging, and the obsessive marketing machine built around it. This is what thickens the fog surrounding AI: not only what the technology can do, but the exaggerated stories, promises, slogans, and misunderstandings constantly piled on top of it.

Let me mention three cases, among the many I have encountered during my professional work with companies on different continents.

I will obviously not name people, companies, or places. The reason, I think, is clear enough.

These cases are included not as gossip, but as practical evidence: AI theatre is not an abstract danger. It appears in meetings, proposals, software purchases, operational decisions, and, sometimes, in expensive mistakes.

The three cases are very useful. Case #1 supports automation versus AI. Case #2 supports using optimisation instead of AI when optimisation is the right tool. Case #3 supports training, accountability, and human verification.

Case #1

A company was modernising part of its software infrastructure. One of the systems under review was a sales management application. As project manager, I had to evaluate several proposals, paying particular attention to their actual effectiveness and to their ability to integrate properly with the company's existing ecosystem.

At one point, one of the managers told me that a certain vendor's solution deserved serious consideration because it "included artificial intelligence". As things stood, however, there was no real need for AI inside that specific sales application. The use of AI had indeed been planned within the broader project, but not directly in that type of software.

So I asked for clarification.

The answer was revealing: "When you enter the postcode, the programme shows you the corresponding area without any possibility of error. So, when the operator enters the customer's address, they cannot get the street name wrong, because the system pulls it from a list linked only to that postcode."

In other words, what was being presented as artificial intelligence was simply a controlled selection based on structured data. Useful, certainly. Sensible, yes. But AI? Not even remotely.

The vendor's marketing department had clearly done an excellent job, especially considering the manager's lack of technical preparation. But calling such a feature artificial intelligence is not technological progress; it is linguistic inflation. AI itself, if it had feelings, would probably feel

insulted at being credited with a task that even the most ordinary relational database can perform without breaking into a sweat.

Detector lesson: A useful database feature does not become AI because someone writes it in a brochure. Before paying for “AI”, ask whether the system is actually learning, classifying, predicting, generating, or interpreting — or simply retrieving structured data.

Case #2

A company needed to assign a number of workers across three different shifts, while respecting a precise set of constraints. Each worker had to be assigned to two shifts, with a defined number of hours between the first and the second. After completing both shifts, each worker was entitled to a full day off. In addition, certain tasks could only be performed by workers with specific qualifications, so the system also had to match people to roles according to their skills and authorisations.

This was not, in itself, a particularly complex software problem to integrate into the company’s ecosystem. It required a sound optimisation approach, not theatrical terminology. A linear programming model, or a well-designed constraint satisfaction algorithm, would have been perfectly adequate. AI could have been avoided here — and, frankly, spared the embarrassment.

However, an AI consultant — or at least someone presenting himself as such — had already been in contact with the CEO, suggesting that AI was obviously the way to go. And indeed it probably was: for his invoice.

Despite the insistence of one of the managers — and insistent managers, as we know, are not an endangered species — the CEO eventually asked me why I had not recommended AI. Given my background, he also asked why I was not proposing to implement it myself.

The answer was simple. This was a mathematical optimisation problem. Linear programming takes the constraints, the boundaries, the available resources, the qualification requirements, the minimum rest intervals between shifts, and the mandatory day off after the two assigned shifts,

then calculates an optimal solution within those limits. It does not need to “imagine” a good schedule. It does not need to “learn” from mystical corporate vapours. It simply solves the problem according to a rigorous mathematical structure.

AI, in this case, would not have improved the solution. At best, it might have approximated a decent schedule. At worst, it could have produced something apparently plausible while accidentally violating one of the hard constraints — assigning a worker to a task for which they were not qualified, placing the two shifts too close together, failing to respect the required rest period, or ignoring the mandatory day off.

The CEO understood the point immediately and dismissed the consultant’s proposal with a courteous but definitive sentence: “We don’t need AI for the sake of AI.”

Fortunately, there are still many intelligent people around the world who can recognise competence where it belongs — and who understand that not every problem becomes more sophisticated simply because someone has managed to attach the word *AI* to it.

Detector lesson: Some problems require mathematics, not mythology. If a task has clear constraints, rules, resources, and objectives, a rigorous optimisation model may be safer, cheaper, and more reliable than adding AI for decorative sophistication.

Case #3

This is a perfect example of why a company cannot simply throw money at technology and hope that technology will politely solve everything.

A company decided to introduce a large enterprise AI tool across the organisation, although it is an SME. It was presented as a major step forward, a productivity revolution, the usual ceremonial fireworks. Management made speeches about “skyrocketing productivity”, but provided virtually no serious staff training. The tool was placed in people’s hands as if the mere presence of advanced software would automatically produce advanced thinking.

Predictably, many employees bought into the hype. They began to treat the system not as an assistant, but as a magic wand. In their minds, it was no longer necessary to check, verify, question, or apply judgement. The machine would do the thinking. They would simply collect the output. There was, of course, one exception: a person who feared losing their job and therefore had the good sense to keep checking what the system produced.

By the third month, the situation had deteriorated. Because nobody had been properly trained, staff began feeding the system complex, ambiguous client requests and dense strategic data sheets, expecting it to make nuanced judgement calls that it was neither qualified nor authorised to make. They were not using AI as a tool within a controlled process. They were using it as a replacement for professional responsibility.

The lack of basic prompt discipline made matters worse. People provided poor context, vague instructions, incomplete information, and sometimes contradictory material. The result was entirely predictable: the AI began producing confident nonsense. It invented data, distorted assumptions, mishandled financial projections, and generated outputs that looked polished enough to pass casual inspection while being dangerously wrong underneath.

But since the staff had already stopped verifying the work, these errors were not caught internally. They moved straight into real client deliverables. Bad assumptions became client documents. Fabricated figures became analysis. Plausible nonsense acquired a logo, a layout, and an air of authority.

The bomb eventually exploded around the sixth month. A client found a serious mistake. Management panicked. The staff were confronted. And then came the most revealing part of the entire disaster: everyone pointed at the computer and said, in effect, "It wasn't me. The AI generated it."

That was the real failure. Not merely a technical failure, but an accountability collapse. The organisation had not only misused the tool; it had allowed people to outsource responsibility to it.

Fixing the situation required much more than tweaking prompts or changing software settings. It required rebuilding the process, reintroducing human verification, defining what AI could and could not be used for, training people properly, clarifying responsibility, and restoring professional judgement where it should never have disappeared in the first place. I can say sincerely that putting things back together was a long and difficult task.

The lesson is simple. Give an untrained team an AI system and tell them it is a magic problem solver, and you do not scale productivity. You scale mistakes, confusion, and irresponsibility at high speed.

Detector lesson: AI does not remove responsibility. Without training, verification, ownership, and clear rules, a company does not scale productivity; it scales confident mistakes.

13 - Notes from the Business Floor

AI is often discussed by people standing very far from the consequences of their own advice.

The loudest voices tend to come from conferences, social platforms, investor decks, consultancy brochures, and people who have recently discovered a new word and are determined to monetise it before the weather changes.

But businesses are not run in keynote slides. They are run in inboxes, meetings, invoices, payroll, product decisions, client expectations, staff constraints, legacy systems, imperfect data, and the daily tension between ambition and cash flow.

For that reason, this book includes short notes from entrepreneurs, founders, consultants, and managers who have had to think about AI from the only place where the question really matters: practice.

Some are enthusiastic. Some are sceptical. Some are cautious. Some have used AI successfully. Some have rejected it deliberately. Some have watched the theatre from the front row and decided not to buy a ticket.

Their role here is not to provide universal truth. No single business experience can do that. Their role is to bring the discussion back to the floor: the place where tools either help, fail, confuse, cost too much, save time, or create yet another system someone has to maintain.

These notes are included because AI decisions should not be made only by listening to vendors, platforms, gurus, or prophets of inevitability. They should also be informed by people who know what it means to build something, serve customers, manage staff, protect margins, and live with the consequences after the sales call ends.

Field Conversation 1

Pete Wilkinson, Founder and CEO, Reclaro

Pete Wilkinson is the Founder and CEO of Reclaro, the Reclaro 1-3-5® OKR Software & Expertise platform. He is also a published author and award-winning speaker. His view is especially useful here because he works inside the software world, yet does not treat AI as something that should be added to a product merely because the market has started shouting.

Pete Wilkinson says:

The first point I'd make is that many companies I come across are still not fully exploring the benefits of AI to their internal and external customers. I think the extent in some companies is using ChatGPT or some similar LLM. But the real benefit is in how AI can streamline operations and remove much of the base level work that takes time.

I do think some companies are afraid or unaware of the benefits AI can bring.

But I do feel people have to be careful about having AI do their thinking for them. I love the phrase that AI is "Assisted Intelligence". Once you understand that it can assist you it gets easier.

When would AI genuinely deserve a place in a SaaS product?

I have some thoughts on this. In Reclaro we have AI built into the product to help people build their plans, set their milestones and then have some basic analysis on how they may reach a goal. But I have to be honest and say it's OK, not brilliant. It certainly helps people when they have a blank screen and don't know what to do next with a goal but it's only a guide.

Where AI could add so much value is in the analytics of the information around goals. We should be able to use AI in a closed environment of course, and have it look at all the goals a company has set and then assist in designing the steps needed to achieve those goals. It should also be able to quickly give valuable Board Meeting level info to assist with the building of board packs for upcoming meetings.

I've come across some companies that have simply added "AI" to the end of their product name. I think that's an easy thing to do but there's a bigger challenge.

That challenge is working out where to add AI into your product that will add genuine real value to the customer not just being a gimmick.

So AI only deserves a place in our product if it helps our customers solve their problems better and faster so they can build better businesses. If it doesn't do that then there's no place for it as far as I'm concerned.

Detector lesson: AI deserves a place in a product only when it helps customers solve real problems better and faster. If it is added merely to decorate the product name, it is not innovation. It is theatre with a feature label.

Field Conversation 2

Luigi Cofone, CEO, Consetra Logistica

Luigi Cofone is the CEO and founder of Consetra Logistica, a logistics, transport, and distribution company operating in the agri-food sector.

Luigi Cofone says:

AI is useful where it is genuinely needed: where it clearly makes work easier, provides accurate information, and increases productivity. If it does not meet these requirements, there is no point in using it.

We have software that does not use AI, yet it is efficient and solves problems very well. We are not going to change those work processes or those types of operations simply to implement an AI application that does not give us certainty. If they already work perfectly, why should we replace them with AI?

Also, although I am not an expert in the field, I have the impression that if AI is fed with incorrect data, it may increase errors and confusion. So we use it only when it is safe to do so.

This is, after all, what you yourself have always taught us.

Detector lesson: If a process already works reliably, replacing it with AI is not progress; it may simply be disruption with a fashionable name. AI deserves consideration when it makes work clearer, safer, faster, or more productive. If it adds uncertainty to a functioning process, the intelligent choice may be to leave the process alone.

Field Conversation 3

Juan Ramón Pérez Berenguer, CEO, Fruklas

Juan Ramón Pérez Berenguer is the CEO of Fruklas, a production and agricultural-services company based in Alicante, Spain. Fruklas operates in the fruit and vegetable sector, with direct grape cultivation and a network of high-quality growers.

Juan Ramón Pérez Berenguer says:

We use business intelligence software that allows us to generate statistics and propose marketing actions to our salespeople. Of course, we feed it with the correct data and verify every piece of information.

We also support one another with production applications that help us with cultivation. In the future, we intend to implement more in this direction. We have some projects, but for us, the use of AI is not an end in itself; it is used only when AI can solve a problem more quickly and reliably.

Detector lesson: AI is strongest when it is connected to verified data, practical operations, and a specific business purpose. It should not be adopted as an ideology but as a tool that helps people act faster, more accurately, and with better information.

Field Conversation 4

Zlatan Tupeković, ISOBELL Division Director

Zlatan Tupeković is the ISOBELL Division Director. His perspective is especially relevant because he uses AI tools directly in his daily work, while remaining attentive to their limits, risks, and wider human implications.

Zlatan Tupeković says:

I mainly use Claude and ChatGPT in my daily work. Claude helps me manage and analyse data, while ChatGPT helps me with small graphic tasks, client email replies, and quick research that I would previously have done manually online. With the right prompt, these tools save time, but I always verify the results.

I see AI as something that should be integrated naturally into everyday work. I am also exploring AI agents for controlled automation of tasks such as research, mailing, and data compilation.

At the same time, there is a real risk that people stop checking what AI produces. If users do not know how to define the right boundaries, context, and references, these tools can generate inaccuracies, misinformation, or poor learning.

AI should not become a substitute for education or knowledge. It should support daily work, especially in areas such as company data analysis, market research, competitor analysis, and, potentially, Digital Twins for testing ideas, products, and market reactions.

I also believe society may eventually need a clear “Made by AI” label to distinguish AI-generated texts, images, news, and content from human-created work.

My deeper concern is human. We already communicate less, and these tools may isolate us further. When my six-year-old daughter asks an AI bot for an answer and assumes it must be correct, I feel my role as a parent is to teach her that knowledge must still be learned, processed, and understood by us.

Human emotion and sensitivity will be very difficult to replace artificially. My concern is that, if this ever happens, it may be because we ourselves have become artificial.

Detector lesson: AI is useful when it supports work, saves time, and helps analysis. It becomes dangerous when people stop verifying, stop learning, or allow the tool to replace judgement rather than assist it.

Field Conversation 5

Jefferson A. Lowe, former director of Planter's Guild Inc.

AI has no personality. It may generate text, organise information, suggest ideas, or accelerate certain tasks, but it does not bring character, judgement, loyalty, humour, intuition, responsibility, or lived experience into a business. Those qualities come from people.

As an entrepreneur, I can benefit from my employees in a way I will never benefit from an AI bot. My workforce is not an interchangeable operating cost to be replaced by whatever fashionable technology appears next. These are people who know the business, understand the customers, recognise the tone of a conversation, notice the small details, and contribute something personal to the company's identity. I will not let them down simply because a machine can imitate efficiency.

This is especially true for small businesses. A small business cannot compete with large corporations by becoming colder, more automated, and more impersonal. Its strength is precisely the opposite: trust, familiarity, human contact, responsiveness, and personality. Many corporations fail in this respect because, as they grow, they often lose their human face. Customers stop dealing with people and start dealing with systems, scripts, forms, ticket numbers, automated replies, and processes designed more for internal convenience than for genuine service.

Of course, a business still needs a good product or service. Personality alone is not enough. But a good product without personality risks becoming anonymous. And when a company starts hiding behind automation, the relationship with the customer can become devious: it looks efficient on

the surface, but underneath it may simply be a way of avoiding responsibility, reducing contact, and making the customer adapt to the system rather than the other way round.

There is also another important point: AI is only useful when you already understand the subject. If you know what you are doing, AI can help you move faster. If you do not, it can become completely useless, or worse, dangerously convincing. It may produce plausible answers, but without human expertise there is no reliable way to judge whether those answers are correct, relevant, or commercially sensible.

For this reason, I do not see AI as a replacement for people. I see it, at most, as a tool in the hands of competent people. The business should not be led by automated systems. It should be led by human beings with knowledge, personality, responsibility, and a real relationship with the customers they serve.

Detector lesson: AI can assist a business, but it cannot replace the human qualities that make a business trusted: judgement, personality, loyalty, responsibility, humour, intuition, and real customer relationships. For small businesses especially, automation should strengthen the human character of the company, not erase it. AI is useful only when competent people remain in charge; without expertise, it can produce convincing nonsense faster than anyone can correct it.

Appendix — The AI Theatre Detector Scorecard

Before approving an AI purchase, score the proposal.

This scorecard is not a scientific instrument. It is a practical judgement tool. Its purpose is to slow the buyer down, force clarity, expose weak claims, and prevent the business from confusing a performance with a solution.

Score each category from 0 to 5.

0 means the proposal is vague, untested, risky, or unsupported.

5 means the proposal is clear, testable, practical, and well controlled.

1. The Claim

Is the claim specific or foggy?

A weak claim says: “This will transform your business with AI.”

A strong claim says: “This will reduce manual invoice data entry by extracting five fields from supplier invoices and sending uncertain cases to human review.”

2. The Problem

Is there a real business pain?

If nobody can name the repeated cost, delay, error, risk, or missed opportunity, the project is probably starting from fashion rather than need.

3. The Process

Is the workflow clear enough to improve?

If the process is not understood, AI may simply accelerate confusion. The business should know what happens now, who does it, where it fails, and what improvement would look like.

4. The Data

Does the system have usable input?

Where is the data? Is it complete? Is it current? Is it accessible? Is it structured enough? Who owns it? Who cleans it? What happens when it is wrong?

5. The Output

Who checks whether the result is right?

The proposal should explain what the system produces, who reviews it, how uncertainty is flagged, and what happens when the output is incomplete, wrong, or risky.

6. The Risk

What happens when it fails?

Failure is not an exception. It is part of the design. A serious proposal includes fallbacks, human review, audit trails, permissions, and safe failure modes.

7. The Cost

Is the benefit measurable?

The buyer should understand upfront cost, monthly cost, usage-based cost, maintenance cost, support cost, and the financial value of the expected improvement.

8. The Exit

Can the company stop using it without chaos?

Can data be exported? Can workflows be documented? Can another provider maintain the system? Are there exit fees? What happens when the contract ends?

9. The Owner

Who maintains the system?

Every AI-assisted process needs ownership. Someone must update documents, prompts, rules, permissions, workflows, and integrations. If nobody owns it, the system will decay.

10. The Proof

Has it worked outside the demo?

A demo is not proof. The proposal should include a pilot, real data, agreed success criteria, and evidence that the tool works under realistic conditions.

Interpreting the Score

A total score of 40–50 suggests the proposal may be worth serious consideration.

A score of 30–39 suggests useful potential, but important questions remain.

A score of 20–29 suggests the project is risky and should not proceed without a clearer pilot.

A score below 20 suggests that the proposal is probably theatre, fog, or a future invoice looking for a host organism.

The scorecard does not replace judgement.

It protects judgement from being chloroformed by fashionable language.

Sienda ltd

www.sienda.co.uk

—

Need to test the problem before choosing the solution?

Visit **Sienda Weblines** — the working hub from Sienda for practical diagnostics, field notes, applied systems, and clearer commercial routes through secure exchange, AI, automation, market intelligence, and business engineering.

www.siendaweblines.com