

THE ELECTE QUARTERLY

NO. 1 — JULY 2026



On the optimal level of artificial intelligence

Contents

THE QUARTERLY ESSAY

How Much AI to Use in Your Business

A guide to the optimal level

01	The AI Laffer Curve	5
02	The Middle-to Middle Principle	7
03	Four Structural Limitations	8
04	The B+ Trap	9
05	A Practical Matrix	10
06	Calibration in Practice	11

Build vs. Buy AI for SMEs in 2026

In-house development vs. outsourcing

01	The AI Imperative in 2026	14
02	What "Build" and "Buy" Really Mean	16
03	Comparative Analysis	19
04	Beyond the Binary Choice	22
05	Making the Right AI Decision	24
06	Calibration in Practice	11

How Much AI to Use in Your Business

A guide to the optimal level in 2026

The most useful answer is not “as much as possible”. It is “to the extent that it adds value without undermining judgment, quality, and differentiation”.

BY FABIO LAURIA — CEO & FOUNDER, ELECTE

This matters more today than it might seem. In Italy, the adoption of artificial intelligence in businesses rose from 8.2% in 2024 to 16.4% in 2024 to 16.4% in 2025, according to Istat.¹ This doubling in just one year tells us one simple thing: the question is no longer whether to take action, but how to adjust the balance.

¹ Istat, adoption of AI among Italian enterprises, 2024-25, reported by *Il Foglio*, December 2025.

As the CEO of an AI platform for European SMEs and as a researcher working on the commoditization of language model outputs, I see the same mistake being repeated. Companies treat AI like a light switch. They either ignore it or try to automate everything. Both choices destroy value. The first because it leaves you lagging behind. The second because it floods you with outputs that are correct on the surface but weak in substance.

The framework that works is simpler and more disciplined: use AI where it can reduce routine tasks, and stop using it where responsibility, context, and human judgment are required.

01 The AI Laffer Curve

Why neither zero nor hundred per cent is the right answer

Most companies make mistakes — either by doing too much or by acting too late. The point isn't to adopt AI. The point is to find the threshold beyond which the increase in operational performance is outweighed by the risk you're introducing. Balaji Srinivasan summed it up better than anyone else:


“0% AI is slow. But 100% AI is slop.”

BALAJI SRINIVASAN

Too little AI leaves the company carrying unnecessary costs. Too much replaces judgment with plausible but interchangeable outputs.

The logic is that of the Laffer Curve applied to knowledge work. At first, every additional unit of AI generates a high return: less time wasted on repetitive tasks, faster execution, and more standardized processes. Then a threshold is reached. Beyond that threshold, the marginal benefit declines and costs begin to rise — costs that many managers fail to recognize until it's too late: well-disguised errors, reduced control, blurred lines of responsibility, and content that all looks the same.

FIG. 1 THE AI LAFFER CURVE



Without AI, performance slows down. With too much AI, you end up standardizing even what should remain distinctive.

Each additional unit of automation returns less once the thresholds crossed; beyond it, hidden costs — disguised errors, reduced control, sameness — overtake the gains.

When Zero AI Is an Operating Cost

Staying at zero isn't prudence. It's choosing to pay qualified people to do work that doesn't create a competitive advantage. It happens every day. Finance teams manually reorganizing files. Sales reps rewriting nearly identical emails. Operations teams moving data from one system to another. Marketing teams preparing first drafts and format changes by hand. These activities don't improve strategy, strengthen positioning, or increase perceived customer value. They simply consume managerial attention and valuable time.

That is why the market is shifting. As noted at the beginning, adoption is growing because the cost of inaction is becoming increasingly apparent — first in terms of time and then in terms of margins.

Use a lot of AI when the work is repetitive, frequent, and easily verifiable.

Scale it back when the output touched money, reputation, trust, or strategic decisions.

When 100% AI Becomes a Slop

The other mistake is more subtle, because at first it seems like a win in terms of efficiency. A financial report written entirely by AI may appear accurate, well-organized, and even convincing. But a responsible CFO doesn't sign off on a document just because it "sounds good." They cross-check it against orders, collections, inventory, operational delays, and business exceptions. Without this step, the company isn't automating effectively. It's merely shifting the risk further down the chain.

Stop it before it signs — before the client report is finalised, before an irreversible decision is made.

The same applies to sales and marketing. An email generated 100% by AI can get the tone, structure, and grammar right. But it often lacks that unique touch: the reference to the customer's actual challenge, the dynamics of their industry, or the specific friction that came up during the call. That's where conversion happens. And that's where total automation starts to erode differentiation.

This is the slop. Content that's readable, quick to produce, and formally acceptable, but lacking in accountability and competitive advantage.

02 The "Middle-to-Middle" Principle and the True Costs of AI

What AI automates well, and where the real cost lie

AI doesn't automate an entire process very well. It automates the core of the process very well. It works "middle-to-middle." At the beginning, you need a human to define the problem, the context, the constraints, and the relevant data. At the end, you need a human to verify the output, put it into context, and take responsibility for it. In between, however, AI can save hours of work.

Take a business analysis, for example. Management defines the initial question: Which customers are seeing a slowdown? Which product lines are growing? Where are margins being squeezed? The AI aggregates data, cleans up tables, identifies patterns, and prepares the report. Then an expert reviews the output and decides whether that pattern is a true anomaly or just temporary noise.

The same pattern applies to customer service, finance, operations, and marketing. AI excels at transformation, classification, summarization, format adaptation, and draft generation. On its own, however, it is not well-suited to setting business priorities or taking on the risk of making the final decision.

Many entrepreneurs focus on APIs or licenses. That's part of the equation, but it's rarely the deciding factor. The real cost lies in the hours of expertise required to provide clear instructions and verify the output.

WHERE THE VALUE SITS

10%	algorithms
20%	data
70%	people, process, culture

THREE IMPLICATIONS

Assign no AI to a process without a human owner.

Never buy the tool first and look for the use case after. Start with the bottleneck.

Measure review time, not generation time alone.

AI does not eliminate the need for expertise. It shifts the focus from mechanical tasks to sound judgment.

03 Four Structural Limitations

Boundaries that show where to stop

The quickest way to get AI adoption wrong is to treat its limitations as temporary problems. Many of them are not. They are structural boundaries that serve precisely to determine where to stop.

i. Costs

Large-scale AI isn't free. Every call, workflow, orchestration, integration, and control adds to the cost. If the task has little value or requires too many review steps, automation can worsen the bottom line instead of improving it.

ii. Mathematical

AI does not magically solve problems where the system is unstable, chaotic, or difficult to observe. A model can help interpret signals. It cannot transform fundamental uncertainty into certainty.

iii. Practical

Even when the model is good, the entire task cannot be fully automated. Someone has to formulate the problem, and someone has to check the answer.

iv. Physical

AI doesn't work on your factory floor, doesn't visit customers, doesn't sense the tension in a negotiation, and doesn't see a machine vibrating abnormally unless someone feeds that information into its data.

68% of Italian firms under 50 employees cite the lack of internal expertise as the main obstacle to adoption; independent use takes 4-6 weeks of training.

THE MANAGER'S FOUR QUESTIONS

Is there any reliable data?

Is there a process owner?

Does anyone know how to validate this?

Is the environment stable enough to make the task repeatable?

If the answer to any of these questions is "no," increase the human quota.

If the process requires implicit context, direct perception, or strong legal liability, AI should act as an assistant, not a pilot.

04 The 'B+ Trap'

How total automation kills differentiation

The most subtle strategic problem isn't a glaring mistake. It's the tendency for good-quality work to converge toward mediocrity. I call this the B+ Trap.

"Good" isn't enough anymore. Major generative models are increasingly producing output that is “good enough.” Clean text. Readable summaries. Well-organized analyses. Correct structures. But when everyone uses the same models, the same prompt patterns, and the same workflows, the results tend to converge.

For many companies, this goes unnoticed at first. They see speed and apparent quality. They don't see the loss of voice, of edge, of competitive advantage. In marketing, this translates to interchangeable content. In analytics, it translates to insights that anyone else can obtain. In strategy, it translates to decisions based on average market intelligence, not on your proprietary advantage.

The advantage lies in the human element. A company that leaves standard tasks to AI and then incorporates internal expertise, industry context, proprietary data, and managerial judgment produces a different kind of output. Not necessarily longer or more complex — but more useful.

This is why 100% AI is a competitive dead end. Not because AI is poor, but because if you let it produce everything without human input, you end up with results that are increasingly similar to everyone else's. The part that creates margin is the non-commodity aspect.

*The advantage in 2026 isn't having access to AI.
It's knowing where to stop automation and add
your own proprietary layer.*

05 A Practical Matrix

The nature of the task, and the cost of error

When an entrepreneur asks me how much AI to use in their company, I start with two factors — not the tool itself. The first is the nature of the task. Is it mechanical, analytical, or decision-making? The second is the cost of the mistake. If the output is wrong, do you lose a few minutes, a customer, profit margin, or credibility?

The most immediate impact of Gen AI is seen in the automation of repetitive tasks such as email management and the generation of standard reports, freeing up human resources for higher-value tasks.

THREE IMPLICATIONS

Corrective intervention rate. If the output requires too many human corrections, you have gone beyond the optimal point.

End-to-end time. If AI reduces production but lengthens the review process, the benefit is modest.

Quality as perceived by the end user. If the customer or team has less confidence in the output, automation has gone too far.

SHARE OF A PROCESS TO DELEGATE TO AI

TASK TYPE	LOW COST OF ERROR	MODERATE COST	HIGH COST
Mechanical, repetitive	~90% <i>Formatting, scheduling, tagging, distrubition.</i>	~70% <i>High automation with final review.</i>	~50% <i>AI drafts; a human reviews line by line.</i>
Analytical, interpretive	~70% <i>AI finds patterns; humans confirm them.</i>	~50% <i>A good balance for management reports.</i>	~40% <i>Systematic expert review required.</i>
Decision-making, strategic	~40% <i>Support for scenarios and options,</i>	~30% <i>AI assists; it does not conclude.</i>	~30% <i>Pricing strategy, hirings, sensitive communications.</i>

These percentages are not a natural law. They are a practical starting point. They help avoid two common mistakes: automating high-risk processes too soon, or leaving processes manual that should by now be automated.

06 Calibration in Practice

The model applied at ELECTE, without theoretical embellishment

The best way to understand this framework is to see it applied without any theoretical embellishments. Internally, the process did not start with an abstract concept regarding the “level of AI.” It began with a simple rule: automate only where the cost of an undetected error is low, while maintaining human control where the cost of error is high.

The clearest example is the editorial pipeline. The first attempt was simple: to automate everything, from the initial draft to distribution across channels, including format adjustments, images, and scheduling. It worked. But the output was only generally correct. The tone was there. So was the format. What was missing was the element that an experienced reader picks up on right away: the specific angle, the judgment, the point of view.

The calibration was achieved by reintroducing human intervention at just two points: reviewing the key message and selecting the angle for each platform. The AI remained responsible for format adaptation, creative production, and publication. As a result, the process was reduced from three hours to about 30 minutes of human work per cycle, resulting in a final balance of approximately 80% AI and 20% human.

The method can be replicated in any SME. Classify processes into three groups — mechanical, analytical, decision-making. Increase automation, then reduce it until quality is acceptable without excessive friction. Set an operating standard and review it every quarter; when any metric deteriorates, the slider moves back.

The sweet spot isn't where the AI can do everything. It's where the team stops over-correcting and the output remains credible.

THE EDITORIAL PIPELINE

3 h → 30 min

human work per publishing cycle,
after calibration

80 / 20

final balance of AI to human

AI should replace the analyst's work when it is repetitive and structured, not business judgment. In other words, it's designed to replace your analyst, not your judgment.

Build vs. Buy AI for SMEs in 2026

In-house development vs. outsourcing

The right choice isn't necessarily the most sophisticated one. It's the one that creates measurable value with the least organizational friction.

BY FABIO LAURIA — CEO & FOUNDER, ELECTE

You're likely facing a very real situation. Your team hears about AI every day, vendors promise efficiency, competitors are starting to make moves, and in the meantime, you have to make a decision that's not just about technology. It's about budgets, priorities, internal expertise, and speed of execution.

For an SME, the question in 2026 is no longer whether to use artificial intelligence. The real question is how to adopt it without creating a project that is costly, slow, and difficult to manage. This is where the dilemma arises: should we develop a solution in-house or purchase a ready-to-use platform?

The choice may seem technical, but it is actually strategic. One approach can offer you more control, while the other offers greater speed. One promises differentiation, while the other reduces complexity and risk. The key is to understand which option delivers real value in your specific context, not in the abstract.

01 The AI Imperative in 2026

Waiting is often the most costly decision in 2026

AI is no longer a lab experiment, nor is it a side project to be put off until the end of the year. It is a decision that affects execution, profit margins, and the ability to react faster than the market. The problem is that the “build vs. buy” dilemma is often oversimplified. “Build” is portrayed as synonymous with control, and “buy” as synonymous with simplicity. In reality, the real difference lies elsewhere: how long it takes you to achieve a useful result, how much risk you’re taking on, and how much complexity you’re introducing into your organization.

Many SME leaders still believe that AI is a priority only for companies with in-house data science teams. That is no longer the case. The pressure stems from very common issues:

- Smaller teams that have to get more done
- Rising costs call for more efficient processes
- Decisions that are made more frequently and require readily available and easily understandable data
- In more volatile markets, forecasting and alerting become operational, not optional

This is the key point that many people overlook. AI in SMEs isn’t growing because it’s “trendy.” It’s growing because it helps manage real-world tasks: automated reports, data preparation, operational summaries, forecasts, and risk management.

AI ADOPTION

35 % UK SMEs were using AI in 2025

25 % increase from the previous year

24 % UK companies plan to adopt by the end of 2026

When a company needs to do more with fewer people, the real benchmark isn’t technical sophistication. It’s the time it takes to turn raw data into useful decisions.

The Cost of Not Choosing

i. Manual processes remain unchanged

The team continues to copy data between spreadsheets, systems, and presentations.

ii. Your organization misses out on learning

While others are testing, making mistakes, and improving, you remain in a passive observation phase.

iii. The market adapts to new standards

If your competitors start responding more quickly to sales signals, forecasting demand more accurately, or monitoring risks more effectively, the gap doesn't stem from an algorithm. It stems from the quality of execution.

Why "Build vs. Buy" is a Strategic Decision

Most mistakes stem from a flawed premise: treating "build vs. buy" as an IT decision.

BUILD VS. BUY DECISION IMPACT

FACTOR	IF YOU TAKE A WRONG TURN
Capital	<i>Commit your budget too early or in a way that lacks flexibility</i>
Timing	<i>Delay the first positive result</i>
People	<i>Overworked, unprepared teams</i>
Governance	<i>A wide range of tools and responsibilities</i>
ROI	<i>It's too late to tell whether AI is actually creating value</i>

For an SME, the key isn't to adopt every possible AI solution. It's about adopting the ones that truly improve work, without turning the initiative into an unmanageable program.

02 What "Build" and "Buy" Really Mean

What separates building from buying isn't technology. It's responsibility.

Many comparisons on this topic are misleading because they rely on overly narrow definitions. “Build” doesn’t simply mean developing a model. “Buy” doesn’t just mean purchasing a subscription. The real choice is about who will shoulder the burden of complexity.

“Build” Meaning

If you choose to build, you’re not just buying freedom. You’re taking on technical and operational responsibilities throughout the entire process.

A BUILD INCLUDES

Data preparation	<i>Collection, cleaning, deduplication, and normalization</i>
Choosing a model	<i>Commercial, open-source, or custom</i>
Integration	<i>Connectivity with ERP, CRM, spreadsheets, databases, and internal workflows</i>
Deployment	<i>Environments, Permissions, Monitoring</i>
Maintenance	<i>Updates, checks, error correction, governance</i>

It's like building a custom-designed facility: more flexibility, but also responsibility for permits, infrastructure, and maintenance.

“Buy” Meaning

When making a purchase, choose a platform or a suite of services that are already designed for common use cases. You’re not giving up on your strategy. You’re simply avoiding building components from scratch that don’t really set you apart.

In practice, "buy" often means:

If your competitive advantage doesn't stem from the model itself, you probably don't need to build the model from scratch.

- preconfigured models
- connectors to distributed data sources
- templates for reporting, forecasting, or alerts
- low-code or no-code interfaces
- maintenance and updates handled by the provider

For an SME, this makes a big difference. The team can focus on processes, KPIs, data quality, and internal adoption, rather than spending energy on architecture and MLOps.

The Intermediate Spectrum

The choice is never entirely black and white. Between building and buying, there are hybrid solutions that many small and medium-sized businesses adopt without even calling them that.

Three common examples:

i. Buy with light customization

Purchase a platform and configure it for workflows, roles, dashboards, and internal data sources.

ii. Buy with API extensions

Use a product that’s ready for common features and add custom components where needed.

iii. Build on purchased components

You don’t start from scratch. Combine APIs, business models, and proprietary logic into a more specialized system.

The most common mistake in SMEs

SMEs often choose to build their own solutions because they fear that buying off-the-shelf software would lead to excessive standardization. But the real question isn't "How customizable is it?". It's "Where do you want to invest your effort?"

If your challenge is automating reporting, forecasting, data preparation, or alerting, the real key to success almost never lies in the model itself. It lies in the operational rules, integrations, and understanding the business context.

On the other hand, if your model or pipeline is a direct part of your competitive advantage, then building it yourself may make sense. But only if you already have a clear understanding of the use case, sufficiently reliable data, and the internal capacity to manage it over time.

Weeks vs. Months

"Buy" — implementation in a few weeks; "build" — 3-6 months

03 Comparative Analysis

The cheapest option is not always the most affordable.

INITIAL ORIENTATION CHART

CRITERION	BUILD	BUY
Initial cost	Taller and less predictable	Spread out over a longer period of time
Time-to-value	Slower	Faster
Required skills	High and continuous	Read more on the inside
Maintenance	To be paid by the in-house team	Managed largely by the supplier
Customization	Top-of-the-line, but expensive	Suitable for standard and customizable use cases
Operational scalability	It depends on the architecture you've created	It depends on the maturity of the chosen platform
Main risk	Delays, complexity, technical debt	Lock-in and adaptation limits

Costs and time-to-value

The first mistake is to focus solely on the upfront cost. The real comparison isn’t CAPEX versus subscription fees. It’s the time and effort required to achieve a result that the business recognizes as valuable.

With custom development, the upfront cost is just the beginning. You need to factor in technical work, coordination, testing, integrations, maintenance, and updates. If the project slows down, costs continue to rise even without delivering operational value. With a buy model, the cost is often more transparent because the provider handles a significant portion of the infrastructure, training from scratch, and model maintenance. This shifts the focus from technical ownership to business outcomes.

For many SMEs, this is a critical factor. If the main constraint is liquidity or the need to deliver results quickly, the predictability of a subscription or usage-based model is more manageable than an open-ended development program.

The problem isn't spending too little. It's spending too late relative to when the business needs the results.

Skills and Maintenance

Building this requires an organization capable of sustaining the AI over the long term. A good developer or a brilliant external consultant isn't enough. You need clear roles, processes, and ownership.

Useful questions are very specific:

- Who prepares and validates the data?
- Who monitors the system's performance over time?
- Who updates the pipelines and models when processes change?
- Who steps up when the business calls for new approaches or new results?

If these answers aren't already clear enough today, the development process risks creating an internal dependency on a few key individuals. For an SME, this vulnerability is often more dangerous than vendor lock-in.

With a buy approach, basic technical maintenance is largely outsourced. This doesn't eliminate internal work, but it does change it. Your team needs to focus on use cases, priorities, data quality, and adoption — not on resolving every infrastructure issue.

Scalability and Risk Assessment

Many people choose builds to “gain control.” But control only makes sense if you can actually exercise it. Having complete architectural freedom is useful when the model, decision-making logic, or pipeline represents a direct competitive advantage. If you're building unique, non-replicable capabilities, this may be the right approach.

For horizontal use cases like search, summarization, and operational support, the AI engine is rarely the differentiator. It lies in the quality of the data, integration with business systems, and governance policies. In these scenarios, purchasing and configuring a solution is often the more rational approach.

BUY VS. BUILD PRACTICAL RISKS

AREA	BUILD	BUY
Execution	Delayed or incomplete project	Vendor lock-in
Evolution	Growing technical and maintenance debt	Limitations on extensive customizations
People	Expertise concentrated in just a few people	Less direct control over the stack and roadmap
Business	Deferred ROI	Risk of choosing an unsuitable platform

This is why the “build vs. buy” debate regarding AI for SMEs in 2026 should be viewed through a managerial lens. The right approach is not necessarily the most theoretically sound one; rather, it is the one that best aligns resources, timelines, and the value that can be achieved.

If your company hasn't yet reached a high level of AI maturity, the biggest risk isn't having less control. It's choosing a level of complexity that you can't manage.

04 Beyond the Binary Choice

Buy to discover. Build to differentiate.

More mature SMEs don't view "build" and "buy" as two opposing approaches. They see them as stages of the same journey.

67% vs. 33%

Success rate of "Buy" vs. "Build" approaches

Buy to Learn, Build to Last

*You buy to learn. Not to become dependent.
You buy to clarify use cases. Not to lock in your
strategy.*

The approach offers three benefits:

i. Shortens the organizational learning curve

The team quickly understands what works, what data is needed, and which processes are truly suitable for automation or predictive support.

ii. Avoid jumping the gun on the wrong customizations

Many companies realize too late that they were trying to build something that a preconfigured platform could have already handled adequately.

Being the first to buy doesn't mean giving up your competitive edge. It means avoiding building in the dark.

iii. Improves the quality of future build decisions

When it comes time to build, you do so with clearer priorities, better data, and more robust operational metrics.

When It's Time to Build

The build comes into play once you've reached a certain level of experience and can confidently answer a few questions:

- Has the use case become central to your competitive advantage?
- Do standard solutions cover the common aspects well but not the distinctive ones?
- Has the team developed enough expertise to manage a custom development project?
- Do you have sufficient evidence to justify adding more complexity?

If the answer is yes, the hybrid model allows you to build only what truly warrants an in-house investment. Everything else is purchased, integrated, or configured.

AI maturity isn't demonstrated by building everything in-house. It's demonstrated by knowing what not to build.

05 Making the Right AI Decision

The best AI strategy is the one your business can sustain.

The choice between building and buying isn't a matter of ideological preference. It comes down to a more practical question: which path will lead your small or medium-sized business most quickly to a profitable, manageable, and sustainable outcome?

Building makes sense when your use case is truly unique and you're prepared to handle the complexity, maintenance, and technical responsibility over time. Buying makes sense when you want to accelerate your impact, reduce internal friction, and keep your team focused on the business — not on infrastructure.

For many SMEs, the most sensible choice in 2026 isn't simply "build or buy." It's to start by buying, learn quickly, validate the value, and build only where it's truly needed. This approach protects budgets, improves time-to-value, and reduces the risk of investing too early in the wrong direction.

If you're making a decision right now, don't go for the solution that looks the most ambitious on paper. Look for the one that makes your company better able to make good decisions, more often, with less friction.

THE ELECTE QUARTERLY



The house journal of ELECTE, published four times a year from Milan

PUBLISHED BY	ELECTE S.R.L. —Via Montenapoleone 8, 20121 Milan (MI)
EDITOR	Fabio Lauria
THIS ISSUE	No. 1 — July 2026 — <i>On the optimal level of artificial intelligence</i>
CORRESPONDENCE	hello@electe.net
COMPOSITION	Set in Playfair Display, New Times Romans and Geist
REGISTRATION	VAT IT 12771670960 — REA: MI-2682918
ISSN	3103-8743