

GUIDE · 2026

AI for CEOs and founders: the method to decide with clarity, and stay ahead

Three original frameworks, two operating matrices, and a 30-day plan to turn AI from a technology question into a measurable competitive advantage.

Contents

Introduction — AI is a strategy question, not a technology question	03
Part 1 — Understanding AI without being technical	05
Part 2 — The 3 pillars of AI competitive advantage	08
Part 3 — What to automate, and how to decide	10
Part 4 — Building governance	13
Part 5 — AI in the manager's daily work	16
Part 6 — The 30-day plan	19
Conclusion — The AI-ready leader	21
About the author	23

INTRODUCTION

AI is a strategy question, not a technology question

In 2026, AI is no longer a question for the future. It is already inside companies, already inside processes, already inside decisions. The question every CEO and founder should be asking is not "when should we adopt AI" but "why aren't we getting real results from what we've already activated".

This is the contradiction I work with every day: companies that have implemented AI tools, bought licenses, run training programs, and still cannot translate any of it into a measurable advantage. They keep operating exactly as before, just with different tools.

The problem is almost always the same: they treated AI as a technology question. They delegated the decision to IT, picked the most heavily marketed tool on the market, and waited for it to "integrate" with existing systems. But AI is not a tool. It is a shift in how decisions get made.

The numbers that matter in 2026

The data leaves little room for optimism for those who stand still:

+40%

productivity increase for knowledge workers who use AI systematically, versus colleagues who don't.

Source: McKinsey Global Institute, 2025

88% / 7%

of companies have activated some form of AI. Only 7% have managed to scale it beyond pilot projects.

Source: Boston Consulting Group, AI Adoption Report, 2025

80%

of business decisions will be supported by real-time AI analysis by 2026.

Source: Gartner, 2024

15–20%

growth gap between companies in the bottom quartile of AI adoption and sector leaders.

Source: Deloitte AI State of the Enterprise, 2025

The gap between organizations that know how to use AI and those that don't is widening at a speed the market has never seen before. It is not a linear phenomenon: it is exponential, because companies that accumulate data and optimize their systems become harder to catch up with every month.

The cost of waiting

There is a pattern I see often in the leaders I work with: "wait and see." The idea that waiting a few months, letting the market settle, watching what competitors do, is the prudent strategy.

It is, in fact, the most expensive choice available.

The reason is simple: AI improves with data. Every day a company delays, the competitor who has already activated an AI system accumulates proprietary data, optimizes its models, lowers operating costs, and decides faster. The gap is not easy to close once it opens.

The good news is that the window has not closed yet. But it is narrowing.

What you will find in this guide

This document is not a technical guide to AI. You will not find system architectures, code, or model comparisons here. It is written for CEOs and founders who want to understand where and how AI can create real value in their business, and who need a method for deciding, not a catalog of tools.

In the pages ahead, you will find original frameworks developed through my advisory work, real cases with verifiable numbers, and a concrete plan you can apply over the next four weeks.

Let's start at the beginning: understanding what is actually happening.

PART 1

Understanding AI without being technical

The 3 waves of automation

To understand where we are today, it helps to put the journey in context. Business automation has developed in three distinct waves, each with a different logic and a different impact on a leader's work.

First wave: data automation (1990s–2010). ERP, CRM, management systems. The goal was to collect data and make it accessible. The value was in order and standardization. The manager of this phase needed to learn to read dashboards and reports. Human judgment remained central to interpretation and decision-making.

Second wave: process automation (2010–2020). RPA, digital workflows, system integration. The goal was to eliminate repetitive, manual tasks: data entry, reconciliations, transferring information between systems. The value was in fewer errors and saved time. But here too, the manager remained the decision-maker. Automation executed; it did not think.

Third wave: decision automation (2020–today). Generative AI, predictive models, autonomous agents. The goal is no longer just to automate execution: it is to augment or replace the decision itself. This is the discontinuity that changes the rules. For the first time, the machine does not simply execute what it is told: it analyzes context, evaluates alternatives, and proposes or makes decisions.

A leader who misses this distinction risks treating AI as a more sophisticated RPA, and missing 90% of the value.

RPA vs. AI agent: the difference that matters

An RPA (Robotic Process Automation) system follows fixed rules. If the data is in field A, copy it to field B. If the invoice exceeds \$10,000, route it to the finance manager for approval. It is fast, precise, and completely rigid: it only works if the world behaves exactly as expected.

An AI agent reasons over uncertain context. It can read a contract written in natural language, identify the relevant clauses, compare them against the company's historical precedents, and produce a risk assessment. It can monitor competitor prices in real time and suggest a repricing action before the sales team even notices. It can analyze customer behavior patterns and trigger a retention action at the exact moment churn probability crosses a critical threshold.

The difference is not one of degree: it is one of kind. And it requires a completely different approach, both in implementation and in governance.

The "paving the cow path" trap

There is a recurring mistake companies make when implementing AI: digitizing existing processes instead of redesigning them.

The term "paving the cow path" comes from urban planning: instead of designing optimal roads, you pave the path cows have already worn over time, because it's familiar and requires less effort. The result is infrastructure that replicates yesterday's inefficiencies with new materials.

In a company, this translates into: "let's take our order-approval process, which takes 7 manual steps and 3 days, and put AI in the middle to make it faster." The result is a marginal improvement to a process that should probably have been abolished or completely redesigned.

AI creates real value when it forces a more fundamental question: "does this process make sense in the age of AI, or does it only exist because we didn't have better tools?" Companies that ask themselves this question find transformation opportunities far more significant than those trying to optimize the status quo.

Why 88% use AI but only 7% scale it

BCG measured one of the most interesting anomalies in the corporate AI landscape: almost everyone has started something, almost no one has managed to take it beyond the pilot stage. The causes are systematic, not random:

Fragmented data

AI models need quality data, in accessible formats, updated in real time. Most companies have data scattered across dozens of different, often incompatible systems, with inconsistent quality.

Absent governance

Who decides what AI can and cannot do? Who monitors its decisions? Without clear governance, pilots stay pilots: the business doesn't trust scaling something it can't control.

Management skill gap

If the leader doesn't understand AI well enough to assess its risks and opportunities, everything gets delegated to IT, which optimizes for the technology, not the business.

Miscalibrated expectations

Pilots get sold internally with promises of extraordinary short-term ROI. When results are good but not extraordinary, the project loses political momentum.

The common denominator among those who succeed in scaling is different: an executive sponsor who understands AI, governance defined before launch, data organized as a strategic asset, and business outcomes, not technology metrics, as the measure of success.

PART 2

The 3 pillars of AI competitive advantage

Not all AI applications create the same kind of competitive advantage. After years of working with companies across different industries, I have identified three pillars that distinguish organizations building a durable advantage from those merely optimizing what already exists without changing their competitive position.

Pillar 1: Learning loops and proprietary data

AI's competitive advantage does not come from access to tools (those are available to everyone) but from the data that feeds those tools.

An AI model trained on a company's proprietary data, continuously updated with feedback from its own decisions, becomes progressively more accurate and harder to replicate. It is an advantage that compounds over time, exactly like the advantage of someone with more experience in a field: more data enters the system, predictions improve, decisions improve, and more data gets generated.

This cycle is called a learning loop, and companies that build one before everyone else accumulate an advantage that becomes exponentially harder to close.

Netflix — recommendation system

Updates in real time with the behavior of 270 million users. About 80% of content watched is discovered through recommendations, not active search.

Source: Netflix Technology Blog, 2024

What this means for your business: the strategic question is not "which AI tool should we buy" but "what proprietary data do we have that no competitor can buy". That data is the perimeter of your AI competitive advantage. Strategy needs to start there.

Pillar 2: Hyper-personalization at scale

AI makes possible something that was impossible just a few years ago: treating every customer, every employee, every supplier as a unique individual, at industrial scale.

Until recently, personalization was a privilege reserved for premium customers, managed manually by dedicated account managers. Today, an AI system can analyze each individual customer's behavior, identify the exact moment they're most receptive to a communication, calibrate the message to their specific profile, and adapt the offer in real time based on their response.

Stitch Fix — personalized styling

AI analyzes sizes, tastes, item-level feedback, and context. 70% of items shipped are kept, versus a 30% e-commerce average.

Source: Stitch Fix Investor Relations, 2024

UPS — route optimization

Real-time route optimization for 100,000+ drivers a day. Cuts 100 million miles a year, saving \$300–400 million.

Source: UPS Annual Report, 2024

Hyper-personalization is not just about marketing. It touches every customer touchpoint with the company, every operational decision, every employee interaction with internal systems. The advantage compounds across every dimension at once.

Pillar 3: Predictive superiority and strategic simulation

The third pillar is perhaps the most underrated, because it concerns the quality of strategic decisions, which is hard to measure in the short term.

Companies that use AI to simulate future scenarios, anticipate market moves, and assess the impact of decisions before making them, build an advantage in the quality of strategic thinking that isn't visible immediately but becomes devastating over time.

JP Morgan — COIN & mosaic intelligence

COIN analyzes loan agreements and flags risky clauses in seconds, replacing 360,000 hours of legal work a year. Mosaic intelligence integrates market data, news, and internal data into predictive signals.

Source: JP Morgan Annual Report, 2024

The key concept here is strategic simulation: the ability to build and systematically test alternative scenarios, with real data, before committing resources. This is what I call Scenario Planning 2.0, and it deserves its own dedicated space in a future installment of this series.

PART 3

What to automate, and how to decide

Understanding AI's potential is necessary, but not sufficient. You need a method for deciding where to start, how to evaluate priorities, and how to choose the right level of autonomy for each application.

In this part, I present two frameworks I use in my advisory work: the AI Process Readiness Matrix and the Agentic Autonomy Matrix. Used in sequence, they give you a precise map of where to act and how.

The AI Process Readiness Matrix

A two-dimensional analysis tool that helps assess which business processes are ready for AI automation and which carry strategic priority.

Vertical axis — process maturity

High when the process is well documented, standardized, executed consistently, and measurable. Low when it depends on tacit know-how, varies in execution, or is hard to measure.

Horizontal axis — potential AI impact

High when the process handles large volumes of structured data, includes repetitive decisions with recognizable patterns, and the improvement is measurable and significant.

QUADRANT 1 · HIGH MATURITY / HIGH IMPACT**Automate now**

Ideal candidates for rapid implementation. Examples: accounting reconciliation, credit scoring, ticket classification, supplier performance analysis.

QUADRANT 2 · LOW MATURITY / HIGH IMPACT**Structure first, then automate**

High potential but not ready. Examples: sales pipeline, onboarding, pricing — standardize before automating.

QUADRANT 3 · HIGH MATURITY / LOW IMPACT**RPA candidates, not AI**

Well-structured, little added value from AI. Examples: e-invoicing, standard reporting, automated notifications.

QUADRANT 4 · LOW MATURITY / LOW IMPACT**Not a priority**

Significant structuring work for a limited return. Push to the bottom of the priority list.

HOW TO USE THE ATTACHED EXCEL TOOL

The file **GMPicks_01_Process_Assessment_v2.0.xlsx** (free download on ceomaze.com) guides you through a systematic assessment of your processes on both dimensions, automatically maps the result onto the matrix, and produces a ranked priority list.

1. List every process you manage (8–15 is the ideal number for a first assessment).
2. For each process, answer the maturity questions (1–5 scale, 6 dimensions).
3. For each process, estimate the potential impact (1–5 scale, 4 dimensions).
4. The tool automatically calculates the positioning and maps your situation.

Spend 90 minutes on this exercise with your team. You'll come out with a prioritized list based on your specific situation, not a generic template.

The Agentic Autonomy Matrix

Once you've identified the candidate processes, this framework helps decide how much autonomy to give AI in each one. Getting this calibration wrong has serious consequences: either systems too rigid to scale, or too much control ceded over high-risk decisions.

● Level 4 — Full autonomy

AI decides and acts without human supervision. Humans monitor aggregate results, not individual decisions.

Appropriate when: very high volume, reversible decisions, limited impact per decision, high demonstrated accuracy. Examples: bid optimization, first-level ticket routing, automatic restocking within thresholds.

● Level 3 — Guided autonomy

AI decides and acts, but logs every action and flags anomalies. Humans intervene by exception.

Appropriate when: high volume, moderate impact, visibility needed without slowing down. Examples: standard customer responses, order approval within thresholds, document classification.

● Level 2 — Human approval

AI analyzes and recommends, the human approves before execution.

Appropriate when: significant impact, hard-to-reverse errors, or a new process without enough data yet. Examples: off-standard contracts, final recruiting decisions, strategic-account pricing.

● **Level 1 — Human decision, AI-supported**

AI gathers, analyzes, and presents the data. The decision is entirely human.

Appropriate when: complex contextual judgment, ethical or reputational implications, human relationships. Examples: layoffs, M&A, crisis management, strategic negotiations.

How to use the two frameworks in sequence

1. **Use the AI Process Readiness Matrix** to identify the 3–5 priority processes for the next 6 months.
2. **For each priority process, apply the Agentic Autonomy Matrix** to define the appropriate level of autonomy.
3. **Build governance based on the level** (covered in the next part).

The 3–dimension impact rule: before deciding the level of autonomy, assess the impact of an error across three dimensions — direct financial impact, customer–relationship impact, reputational impact. If even one carries high impact, don't go beyond level 2 in the initial phase.

PART 4

Building governance

This is the chapter most AI guides skip, and it's the one that makes all the difference between a successful pilot and a scalable implementation.

AI governance is not a document. It is an operating system that defines who decides what, how systems are monitored, and how you adapt when something doesn't go as planned.

Guardrails: where AI's autonomy ends

A guardrail is an explicit constraint that defines the boundaries within which an AI system can operate autonomously. Without clear guardrails, a system either gets limited unnecessarily out of fear (losing 90% of the value) or operates without control (exposing the company to unmanaged risk).

Boundary conditions

Situations in which the system must always stop and hand off to a human, regardless of its confidence level. Absolute rules, not probabilistic thresholds.

- "No credit decision above \$50,000 without human approval"
- "No communication with customers in active legal dispute"
- "No price-list changes beyond $\pm 15\%$ of base price"
- "No action on accounts with revenue above \$X without confirmation from the account manager"

Boundary conditions are defined before activating any AI system, in collaboration between the process owner, legal, and, where relevant, finance.

Confidence threshold

Most AI systems produce a prediction with a confidence level (e.g., "this transaction is likely fraudulent with 94% confidence"). The confidence threshold defines the value below which the system must stop acting autonomously and flag a human.

The right threshold varies enormously by process and consequences. For ticket routing, 70% may be acceptable. For approving a loan, probably not. It needs to be monitored over time and adjusted based on real-world performance.

Anomaly triggers

Signals indicating something unexpected is happening, even if the AI system hasn't detected a specific error. When one fires, the system suspends its autonomy and alerts a supervisor.

- A 30% increase in request volume versus the 4-week average
- A drop in the acceptance rate of AI recommendations below a critical threshold
- A significant gap between model predictions and actual outcomes over a recent period

Monitoring: what to measure, and how often

An AI system without monitoring is a system out of control. Monitoring is a responsibility of the process owner, not just a technical activity.

Accuracy metrics

How often the system makes the right call, measured against real outcomes.
High-volume processes: monitor daily.
Low-volume: weekly is enough.

Business outcome metrics

Is AI producing the expected business results? If accuracy is high but outcomes don't improve, the problem is in the action that follows, not the model.

Operational trust metrics

How often do human operators override AI recommendations? A high override rate signals low accuracy, an adoption problem, or counterintuitive (and possibly correct) recommendations. Review at least monthly.

The continuous learning loop

An AI system that doesn't learn from its mistakes degrades over time, because context changes while the model stays fixed. Three mechanisms keep it current:

Explicit feedback: every correction from a human operator enters the training dataset as a signal. The system learns from its mistakes systematically.

Implicit feedback: end-user behaviors (a customer who ignores a retention email, or opens it but doesn't buy) are signals the system can use to update its predictions.

Periodic retraining: independent of feedback, the model needs retraining with the most recent data. Stable industries can retrain quarterly; high-volatility industries (pricing, demand) need to do it more often.

3 real cases with numbers

JP Morgan — Contract Intelligence

360,000 hours/year of manual contract review replaced by seconds-long analysis, lower error rate. Estimated savings exceed \$150 million. Level 2 autonomy: weekly override review, quarterly retraining.

Source: JP Morgan Annual Report, 2024

Klarna — Customer Service AI

67% of requests handled autonomously in month one (levels 3–4). Response times dropped from 11 to 2 minutes. Guardrails: automatic escalation for legal complaints and large refunds.

Source: Klarna Press Release, 2024

Rolls-Royce — Predictive Maintenance

4,000 engine parameters monitored in real time. Unplanned downtime cut 25%, maintenance costs down 15–20%. Level 2 autonomy: human engineer validates every intervention.

Source: Rolls-Royce Holdings Annual Report, 2024

PART 5

AI in the manager's daily work

So far we've covered frameworks and corporate cases. This part is different: it's more operational, more personal, and it's about how your daily work as a leader changes when you integrate AI systematically.

The 3 applications that change your day

1. Faster decision preparation

One of the activities that consumes the most time in a leader's job is gathering and synthesizing the information needed to make a decision. A well-configured AI system can compress this time dramatically: instead of requesting a weekly report from your analytics team (2–3 days to produce), you can ask directly: "give me the 5 most relevant signals on my key accounts from the last 7 days, with a recommended action for each." Two minutes, instead of hours or days.

2. Scenario simulation before deciding

With AI, you can quickly build alternative scenarios, test your assumptions, and identify the critical variables that could change the outcome. This doesn't replace judgment: it enriches it — the same difference between a pilot flying by sight and one with full instrumentation. Both can fly, but in low-visibility conditions, the instrumentation makes the difference.

3. Transforming meetings

Meetings are the least efficient format for sharing information. AI can eliminate the "bringing everyone up to speed" part: if participants arrive already briefed by personalized AI dashboards, the time together can be spent entirely on reasoning and deciding. Some companies already generate the meeting "pre-read" automatically, tailored to each participant's role.

Turning data into decisions: the "company brain"

An integrated system that aggregates every relevant data source for your business, synthesizes it into actionable insight, and makes it available in real time. Mid-sized companies are building this today with accessible tools. The basic architecture includes:

- **Structured data:** CRM, ERP, financial data, operational metrics
- **Unstructured data:** sales emails, customer feedback, call transcripts, meeting notes
- **External data:** market trends, competitor moves, demand signals
- **AI layer:** a model that integrates these sources, identifies patterns, and answers questions in natural language

Instead of opening 6 different tools, you have a single point of access that answers: "What is the churn risk for my top 10 accounts next quarter?" or "Which deals are at risk this week, and why?"

10 operational prompts for leaders

Designed for any advanced LLM (GPT-4o, Claude, Gemini). They require you to provide the specific context of your situation — structured questions designed to extract strategic value from a conversation with AI.

PROMPT 1 — RAPID COMPETITIVE ANALYSIS

You are a senior strategy analyst. I need a rapid analysis of [company name]'s competitive position in the [specify] segment. Analyze: (1) the main points of differentiation versus competitors [list them], (2) the most exposed vulnerabilities, (3) the 3 moves a competitor could make in the next 6 months to attack us. Be direct and avoid obvious observations.

PROMPT 2 — STRUCTURING A COMPLEX DECISION

I need to make a decision about [describe the decision]. The options I'm considering are [list them]. Help me structure the decision by identifying: (1) which assumptions I'm making that could be wrong, (2) what information I'm missing that would change my assessment, (3) the criteria I should use to evaluate the options, and (4) the asymmetric risks for each option.

PROMPT 3 — PROJECT PRE-MORTEM

Our team is about to launch [describe the project]. Imagine we are 12 months in the future and the project has failed significantly. List the 8 most likely causes of failure, ranked by probability. For each one, suggest a preventive action we could take now. Be specific and don't limit yourself to the obvious causes.

PROMPT 4 — SYNTHESIZING WEAK SIGNALS

Here are 5 recent observations from my market [describe them, even informally]. What do these observations have in common? What trend or shift might they signal? What would be the main strategic implication for a company like ours [briefly describe the business]?

PROMPT 5 — STRESS-TESTING A STRATEGY

Our strategy for the next 12 months is [describe in 3–5 points]. Play the role of a skeptical consultant who has seen many similar strategies fail. Identify: (1) the 3 most fragile assumptions, (2) the market scenarios where this strategy fails completely, (3) what you'd change with a budget 30% lower than planned.

PROMPT 6 — CHURN RISK ANALYSIS

Here is data on a customer [or segment]: [insert relevant data]. Assess churn risk on a 1–10 scale and justify the assessment. Suggest 3 specific retention actions to activate in the next 2 weeks, ranked by impact-to-effort ratio.

PROMPT 7 — NEGOTIATION BRIEFING

I need to negotiate [describe the subject] with [describe the counterparty]. My goal is [specify]. Their current position is [describe]. Help me prepare: (1) my 3 strongest arguments, (2) concessions I can make without compromising the key outcome, (3) likely objections and responses, (4) my BATNA.

PROMPT 8 — DIAGNOSING A PERFORMANCE PROBLEM

One of my key metrics [specify] has worsened by [specify %] over the last [weeks/months]. The context is [describe]. Help me build a tree of possible causes, starting with the most likely. For each, suggest the data point that would confirm or rule it out.

PROMPT 9 — BUILDING A FUTURE SCENARIO

In my industry [describe it], I'm trying to understand how the situation might evolve over the next 18 months. The variables I consider most uncertain are [list them]. Build 3 distinct scenarios: optimistic, base case, pessimistic. For each, describe our relative competitive position and the 2 priority actions to take.

PROMPT 10 — PLANNING AN ORGANIZATIONAL CHANGE

We need to implement [describe the change] in our organization. It involves [number] people and affects [describe the processes]. Help me design a 3-phase change management plan that minimizes resistance and maximizes adoption speed. For each phase, specify: objective, key activities, duration, and how to measure success.

PART 6

The 30-day plan

Everything covered so far is only useful if it turns into action. This part gives you a concrete, week-by-week plan to move from theory to your first working AI pilot in 30 days.

WEEK 1

Map your processes

Day 1–2 List every recurring process in your organization that involves repeated decisions or data analysis. Don't filter yet — write down everything.

Day 3–4 For each process, assess maturity (how standardized it is) and potential impact (time saved, error reduction, revenue effect) using the AI Process Readiness Matrix from Part 3.

Day 5 Build your process map and select your top 3 candidates for automation, ranked by impact-to-effort ratio.

Expected output: a prioritized list of 3 processes, each scored on maturity and impact, with a clear top candidate.

WEEK 2

Pick your first candidate and design it

Day 1–2 Document the current process in detail: who does what, what data is used, what decisions are made, and what the current error/exception rate looks like.

Day 3 Define the target autonomy level using the Agentic Autonomy Matrix from Part 3, and draft a one-page AI Governance Brief: guardrails, confidence threshold, anomaly triggers.

Day 4–5 Choose the tool or vendor for the pilot, and design the pilot's scope: which subset of cases, volume, and duration.

Expected output: a one-page AI Governance Brief and a defined pilot scope ready to build.

WEEK 3**Build the guardrails and launch the pilot**

Day 1–2 Technical setup: connect data sources, configure the chosen tool, implement the guardrails defined in Week 2.

Day 3 Brief the team that will work alongside the system: what it does, what it doesn't do, when to override it, how to give feedback.

Day 4–5 Go live with the pilot at limited scope, with active human monitoring on every decision.

Expected output: a live pilot, a briefed team, and a monitoring routine in place from day one.

WEEK 4**First pilot in production, and evaluation**

Day 1–3 Collect data on accuracy, business outcomes, and override rate, per the monitoring framework from Part 4.

Day 4 Run a team retrospective: what worked, what didn't, what surprised you.

Day 5 Decide next steps: scale the autonomy level, expand the scope, or go back and redesign.

Expected output: a documented pilot result and a clear decision on what comes next.

CONCLUSION

The AI-ready leader

Where am I making decisions today based on intuition and experience that could be significantly improved with more data and systematic simulation?

This is the question this guide has tried to help you answer, not in the abstract, but with frameworks, real cases, and a concrete plan you can start using this week.

The skills that matter over the next 3 years

As AI takes over more structured tasks, the skills that differentiate a leader shift. Four stand out:

Critical thinking about data

The ability to question a model's output, spot when a confident answer is wrong, and know which assumptions sit underneath every prediction.

The ability to ask quality questions

AI is only as useful as the question you bring to it. Precise, well-framed questions produce dramatically better answers than vague ones.

Governance and risk judgment

Deciding where AI should and shouldn't have autonomy is a judgment call that requires business context no model has.

Speed of learning

The specific tools will change. The leaders who win are the ones who adapt fastest to whatever comes next, not the ones who master today's tools.

The question worth asking yourself

"Am I building an organization that gets better every week because it learns systematically, or one that simply repeats what it has always done, a little faster?"


What's next

This guide focused on the foundation: frameworks for deciding where and how to use AI, with the discipline of governance. A second volume, dedicated entirely to agentic AI, the systems that don't just suggest but execute autonomous chains of actions, is in preparation, for leaders who want to go one level deeper.

If you'd like to discuss how these frameworks apply to your specific organization, you can find advisory and consulting services through Giacomo Melani's work at [**giacomomelani.com**](https://giacomomelani.com).

ABOUT THE AUTHOR

Giacomo Melani



Giacomo Melani has more than 20 years of experience in strategy, innovation, and technology, working across industries from telecommunications to manufacturing to professional services. He has led growth initiatives for companies ranging from early-stage scaleups to established multinational groups, and holds two registered patents.

He founded CeoMaze to bring these frameworks to an international audience of CEOs and founders, and writes on giacomomelani.com for the Italian-speaking market.

€600K→€2.5M

scaleup revenue growth in 3
years

+30%

sales increase,
Vodafone
partnership

2

registered
patents

20+

years of
experience

CEOMAZE

ceomaze.com | info@ceomaze.com

© 2026 CeoMaze. All rights reserved. This document is intended for the personal and professional use of the reader. Reproduction, in whole or in part, without written authorization is prohibited.