7 ORIGINAL SINS OF AI INVESTMENT

Entermind

THE 7 SINS

PROLOGUE

> **More than 800 million people use ChatGPT every week…AI has gone from something people play with to something people build with every day.[1]**
>
> **Sam Altman, co-founder and CEO of OpenAI**

Each week brings a new manifestation of the miracle we call Generative AI. The first technology in human history that can speak (or say, autocomplete) our natural language, triangulate decisions, tap into memory and even "reflect". It's still raw and backprop-based, far from how humans actually think, but miraculous nevertheless. Enough to make the bazaar go bonkers with shrill sirens of singularity.

Hype, however, has its place by serving a civilisational purpose. It focuses humanity's resources where it can push the limits of what's possible. It draws in the sharpest minds like Altman's 10,000x researchers whose work can transform entire industries. It unlocks unprecedented capital like Mark Zuckerberg's billion-dollar superintelligence pay packages. And it fuels a global FOMO powered by the futurati: the media, VCs, consultants, influencers and evangelists. All urging the world to move fast, break things and get on board before it's too late.

All of this is fine, until it isn't. There is a thin line between a martyr and a gamechanger. Most companies are not in the business of AI; they are in the business of serving their customers. AI can help them operate cheaper, faster and better, but it still wouldn't be their core business. So, they must undertake the essential task of deciding which AI capabilities are genuinely useful to them. In the end, AI must create customer value, enhance it or help deliver it more efficiently. Everything else is secondary.

A rather revealing MIT report found that about 95% of Generative AI pilots failed to deliver measurable business impact.[2] To put that in perspective, the failure rate of tech startups has reached 92% in 2024.[3] It's astonishing to note that early GenAI investments carried a risk profile similar to early-stage startups, often without the upside of becoming a unicorn. That is an extraordinary level of investment with little to no returns. Not every business leader or CXO has the appetite for that kind of risk.

CXOs need to step back and take a long hard look at how AI is being used and where it's leading them. In this whitepaper, we build on our experience overseeing 700+ AI projects across Asia-Pacific, Latin America, Middle East and North Africa, along with insights from other valuable sources often overshadowed by hype in the AI space. Here, we outline the 7 most common sins behind first-generation enterprise AI failures — and what businesses should be doing instead.

Tech Startup Failure Rate (2024)

Source: AI4SP, "Why 90 % of AI Startups Fail"

FORCING AI INTO LEGACY WORKFLOWS

Silicon Valley loves to talk about AI workers — agentic or otherwise. These digital workers are meant to work alongside human employees, which means that if AI agents or applications are ever to scale, we must learn to industrialise human-AI teaming. Agents can be woven into workflows in many configurations. They can fully replace a human, or they can assist to increase human productivity in a way that is faster and better. In some cases, they can even produce work outputs that is otherwise not possible without AI.

Humans in turn, may guide the AI agent, validate it and handle outlier cases. What you cannot do is take an AI tool and force-fit it onto an existing workflow as it is bound to create a bottleneck down the line.

SIN 01

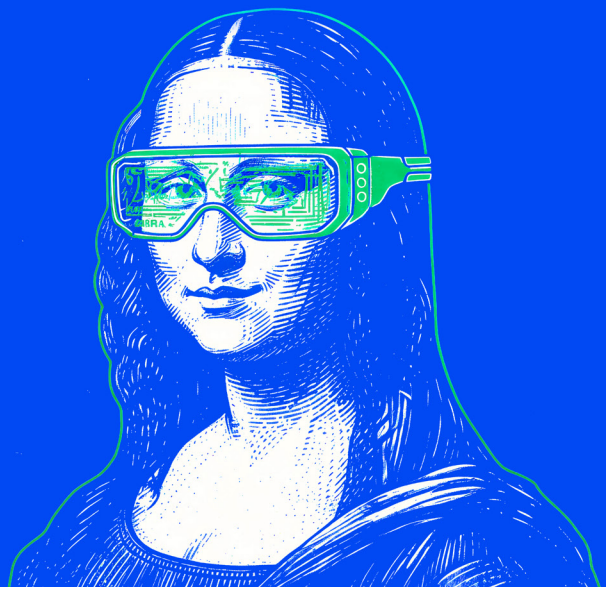

You need an AI-native workflow.

Workflows should be redesigned from a blank sheet, combining the unique capabilities of humans with distinctive strengths of AI. On the behavioural side of human-AI teaming dynamics, certain unique psychological aspects come into play.

In online games, players often prefer human teammates over NPCs.[4] Similar patterns can emerge with AI co-workers; people don't feel recognised when selected by an AI and they feel less guilt about taking credit for AI-generated contributions.

Alpha users rely on AI to speed up tasks, yet much of the productivity gains remain invisible to organisations.[5] Employees still want the ability to override AI recommendations. When that sense of agency slips, anxiety and loss of morale tend to follow. Even roles designed with human-in-the-loop to validate AI decisions can erode as people slowly become passively dependent.

The fear is not that AI starts behaving more like humans, but that humans start behaving more like AI.

At this interface of humans and AI, there is both deep insecurity as well as curious excitement; a struggle to retain agency alongside creeping passivity; empowerment mixed with disconnection. Yet, somewhere in that tension lies the ideal human-AI teaming that can yield a multiplier effect on productivity.

Humans thrive at work when three fundamental psychological needs are met, as established in Deci & Ryan's Self-Determination Theory (1985).[6]

1 AUTONOMY

Let humans steer, not serve AI systems.

Decision Latitude	: Choose how to approach tasks or clients
Method Flexibility	: Ability to decide how to perform a task
Voice & Influence	: Confidence that one's input shapes key decisions
Goal Ownership	: Working toward meaningful outcomes, not just compliance
Values Alignment	: Acting according to personal ethics, not only profit

2 COMPETENCE

Design AI that augments skill and gives feedback loops.

Skill Mastery	: Desire to build expertise
Feedback & Clarity	: Knowing how well one is performing
Challenge-to-skill Fit	: Work that is neither boring nor overwhelming
Professional Growth	: Visible career progression & new responsibilities
Innovation Confidence	: Belief that experimenting is safe

3 RELATEDNESS

Position AI as a trusted partner helping them fulfil meaningful goals.

Team Belonging	: Feeling part of a trusted unit
Recognition & Respect	: Appreciation from peers & leadership
Client Connection	: Building an authentic bond with customers
Organisational Identification	: Pride in the institution's purpose
Psychological Safety	: Freedom to question, dissent or admit mistakes

How AI integration affects these three need states; whether it is additive, subtractive or a mutually beneficial win-win value exchange will determine if humans will allow an AI-native workflow to scale.

BUILDING FOR DEMOS, NOT FOR THE FLOOR

Too many first-generation AI applications were built with a lift-and-shift mindset, focused on what the technology could do rather than what everyday processes require. The messy practicalities of business were treated as background noise, oversimplifying pilots. When scaled, these pilots either broke or hit constant roadblocks, frustrating users. Even when functional, the value delivered was often too underwhelming to justify behavioural changes, hence the tools were quickly abandoned.

AI must start with the nuance of the floor, not as an afterthought. Bottom-up, one workflow at a time.

SIN 02

The workflow needs to mirror real processes.

The first generation of foundation models treated compute as the moat. But as billions poured in and compute edge became less relevant, data became the new moat — especially proprietary data accessible to the model. Today, models are pushing further, attempting to distil expertise and best-practice workflows from top practitioners in the most capable organisations.

Yet broad dataset training naturally pulls models toward the statistical middle, producing safe, general and average answers. Agents built on these models inherit the same workflow assumptions and default to the path of least resistance, thus regressing to the average. That's why the next real moat will not be data, but deviation. How unique, nuanced and high context a business' workflows are, will become its unique advantage.

Failing to capture that in enterprise AI applications means sacrificing the one edge against the homogenising law of averages.

While much of the AI debate focuses on models and agents, success in real businesses like a retailer, bank, telco or FMCG company is often 90% orchestration and 10% technology.

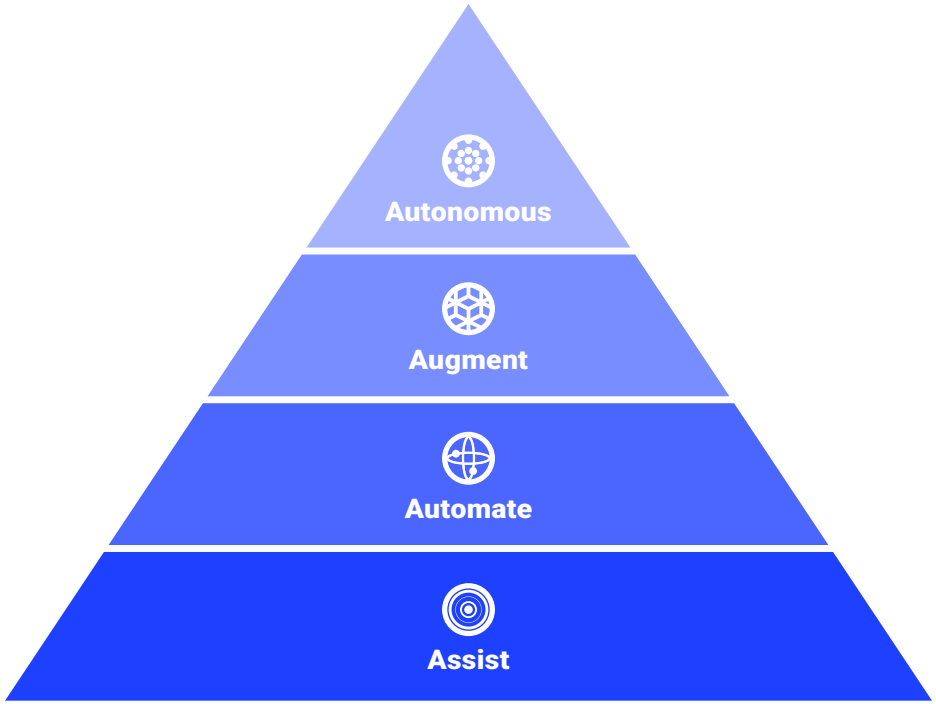

DESIGNING STATIC AI THAT DOESN'T LEARN FROM HUMANS

The open web is humanity's closest form of living memory, even if it represents only a small sliver of the total (public and private) knowledge sphere. Most AI models are built from it. And while generative AI is delightful for general use, it often struggles with narrow, specialised workplace applications.

Enterprise AI compounds this problem. By the time a system is built and deployed, tools like ChatGPT may have evolved so quickly that they outperform it, despite all the fine-tuning invested in the enterprise app. Overengineering and excessive guardrails can further limit the model's capabilities, stunting its natural power and fluency. However, the deeper issue lies in poor data integration and missing feedback loops.

SIN 03

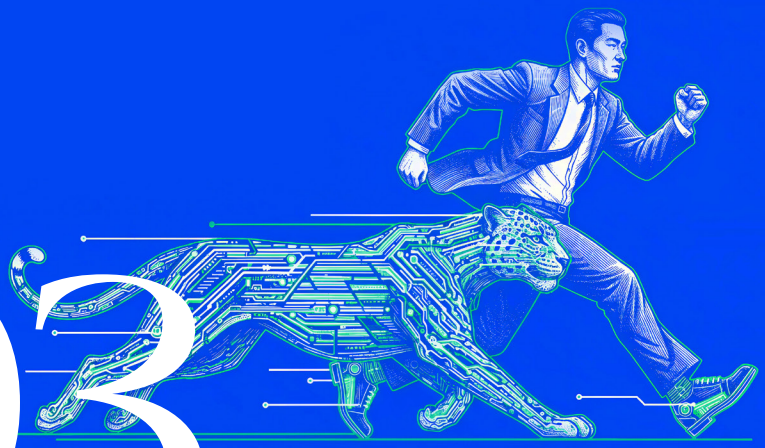

AI must learn from humans and get better with every use.

RLHF (Reinforcement Learning from Human Feedback) is probably the single most powerful process by which AI becomes smarter on context and personal preferences. It creates an exponential user network effect that grows into a sustainable competitive advantage. Just as RLHF is great for diversity tuning, RAG (Retrieval-Augmented Generation) is great for depth tuning by drilling deep context into the application. Blending first-party data (your own), second-party data (suppliers or publishers), and third-party data (off-the-shelf paid data) with a dash of sensible synthetic data can create outputs and decisions that are far more rigorous.

Prompt assembly and guardrails can then be continuously refined to align the technology with the true levers of business effectiveness.

> **Approximately 80% of enterprise data is unstructured…Less than 1% of this data is in a format suitable directly for AI consumption.**[7]
>
> - IBM

Source: Workflow of automated RLHF (Reinforcement Learning from Human Feedback)[8]

Many underestimate how critical it is to prepare this data into an AI-ready state to enable minimal latency, optimal transparency and maximum value. Fortunately, a new generation of AI-powered tools is emerging to ease the burden of data wrangling. How data is architected determines which use cases can be built and which customer experiences can be delivered. Getting it right is essential to unlock the true learning potential of the AI application that sits atop the stack.

Now that machines can speak with fair eloquence, it's important to realise that communication is more than just words. It is tone and manner, dialect and nuance, quantity and quality, and the ability to convey empathy with a pause here and an exclamation there.

For people using AI in their everyday lives, accuracy is only half the game; chemistry is the other. To train AI for this adaptive chemistry, it needs not only left-brain tuning for relevance, but also right-brain tuning for relationship. Employees and customers must feel comfortable opening up, collaborating and helping it learn. AI must not only be correct; it must also connect.

Empathy is the true gift of generative AI. Now that the technology can talk to people, it must learn to "speak" to them too.

IF YOU OWN THE HOUSE, YOU ARE OWNED BY THE HOUSE

As talks of disruption intensify across every category, many companies have rushed to set up their own internal AI units by hiring data teams and launching a flurry of AI apps and agents. While this certainly helps to raise the organisation's overall AI quotient, a recent MIT-affiliated report found that external partnerships for AI development achieve significantly higher success rates, roughly twice that of in-house builds.

SIN 04

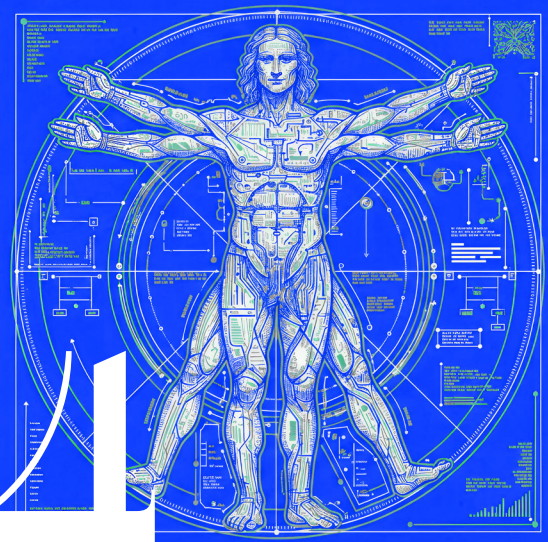

Find a balance with a hybrid model of internal teams and external vendors.

Building a robust AI application requires a multiplicity of skills that vibe coding alone cannot provide.

In-house teams may lack depth across ML engineering, data architecture, UX, orchestration, evaluation and rapid iteration, especially in a technology landscape that evolves every quarter. They may also lack the agility, optionality and competitive pressure that specialised external vendors bring. Internal cost structures can become heavy. Over time, in-house teams may have less incentive to outperform external partners who must stay competitive to survive.

This is one example of how an AI agenda can quietly become captive to the team that was hired to liberate it. But it doesn't mean an internal team is always a bad idea. Some roles fit best internally, while others are better handled by specialised partners. A judicious balance between in-house and outsourced teams is essential to capture both speed and cost efficiency for enterprises.

Success Rates for Enterprise AI Buys vs. Builds

Strategic Partnerships (Buy)

66%

Procure external tools and co-develop with vendors

Internal Development (Build)

33%

Build and maintain GenAI tools fully in-house

Source: MIT NANDA, State of AI in Business 2025, The GenAI Divide[2]

OPERATING AS CENTRALISED ISLANDS INSTEAD OF FEDERATED RIVERS

Too many AI implementations have been tightly centralised, orchestrated by a small core team. While this may sometimes accelerate rollout, it often fails to get real buy-in across departments and business units.

Pilots may look great on paper but frequently fall flat in practice. Even when it scales, costs often balloon because the solution wasn't designed to incorporate real-world capabilities and acclimatisation needs. The challenge then, is how to orchestrate AI in a decentralised, bottom-up way while maintaining speed, executive accountability and some level of standardisation?

SIN 05

You need to operate AI workflows as a river, not an island.

Effective AI orchestration must draw on the classic playbooks of change management. Alpha users already exist within every team, leveraging tools like ChatGPT, Claude, and others to enhance their day-to-day work. They understand workflows and know exactly where AI can add value.

By harnessing their curiosity to drive adoption while pairing them with centrally appointed AI engineers who understand the enterprise tech stack to close last-mile gaps, AI can become a practical everyday default. This approach also builds broad-based credibility and accelerates organisation-wide adoption.

Change Management for AI Integration

	Pair users with AI engineers		Monitor progress		Improve with constant feedback loops
1	2	3	4	5	6
Identify alpha users		Leverage data & workflow knowledge		Address blockers	

In an ideal orchestration, AI is not a mysterious black box locked in the CIO's office, where everyone must bow to at the expense of their priorities, KPIs and sanity. Instead, it is a federated mesh, with individual teams managing not only their data quality, consent and ingestion cadence, but also their models, guardrails, agentic configurations, token budgets, and AI awareness imperatives. All of this is loosely bound by an AI constitution that defines enterprise-wide principles.

How data flows within an organisation reveals much about its culture. When harvested, architected, processed and deployed effectively, data can make a business transparent, resilient, prescient and adaptive. The right data culture can transform a slow, siloed, top-down company into a living, shapeshifting titan powered by live knowledge systems that turn dormant data into actionable intelligence. Systems that learn from user networks can evolve to deliver sharper decisions, superior customer value and persistent high margins.

Data can truly be the beating heart of change.

A highly centralised approach to data and AI orchestration not only increases the risk of scaling failure; it also misses the larger opportunity to raise the bar for the entire organisation. AI should not be viewed as merely a technology initiative, but also a transformation opportunity.

PRIORITISING TECH METRICS OVER BUSINESS IMPACT

SIN 06

Eventually, even a miracle must become an accounting line. Silicon Valley future makers may afford long returns horizons, but the usual listed companies are answerable to shareholders every quarter. That means any significant AI investment must show up either in revenue growth, increased profitability, improved customer satisfaction, stronger compliance or other measurable business outcomes.

This requires enterprise AI rollouts to embed the value case at the design stage itself.

AI programs must be designed for P&L first, not tech first.

Each application should be evaluated for its potential business impact:

- What staffing productivity gains are expected?
- What cost savings can be realised?
- How much additional customer acquisition, upsell or cross-sell can AI enable?
- What uplift in NPS or customer experience is plausible?

These questions keep AI initiatives grounded in business reality rather than drifting toward technical advancements alone. A value-first approach also ensures a balanced AI portfolio of programs that can unlock new growth, reduce staffing and operational costs or improve product and service quality in ways that compound customer value.

As pilots scale, organisations should establish an AI Value Office to track, manage and optimise these levers. An AI P&L process with allocated token budgets and productivity targets helps maintain discipline, while smart budgeting ensures dynamic optimisation as real-world learnings accumulate. Ultimately, AI should be designed to serve the business, not the other way around.

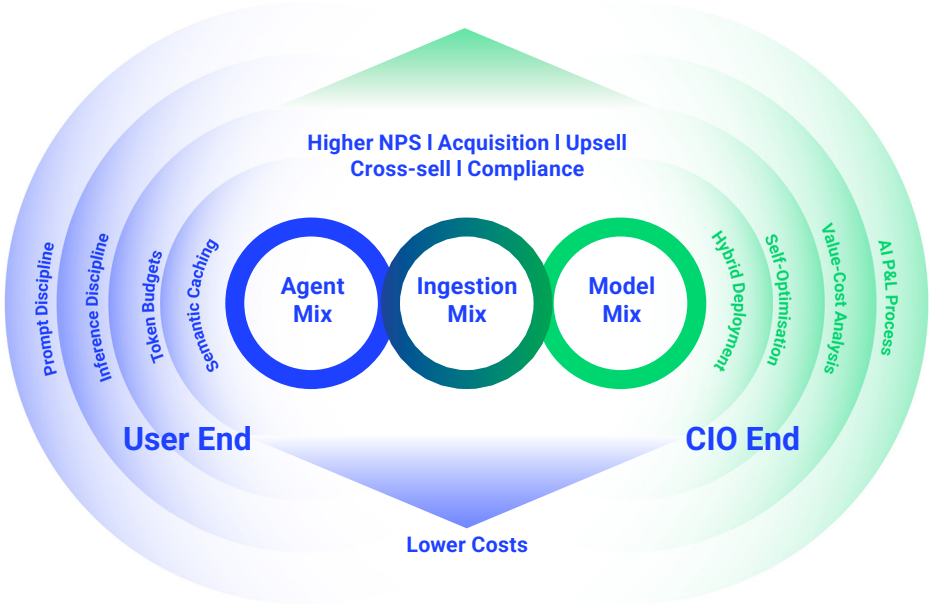

TRAINING BRAINS YOU DON'T OWN

Every time you use an AI application, you are training it; often in ways that go far beyond simple prompts. Chances are, it might have digested proprietary data and developed a deep understanding of your nuanced workflows. It will recognise what works and what doesn't for your business over time.

That accumulated learning becomes an essential part of your company's core intelligence layer.

SIN 07

You need to own the brain you train.

Failing to own and ensuring the intelligence is portable across platforms means you risk losing a critical advantage to either the external platform or competitors (or both).

More importantly, this is not only about proprietary data leaving your secure environment. Even if your data never moves an inch, the learnings generated from daily use via agents and workflows may still circulate back into a shared model, gradually being redistributed across the industry. Bit by bit, what makes you unique gets diluted.

To avoid this, know exactly where your intelligence lives, which parts you own and what happens to it if you decide to part ways with your tech vendor one day. Remember: trade freely, but do not trade away your freedom.

Layers of Business Intelligence

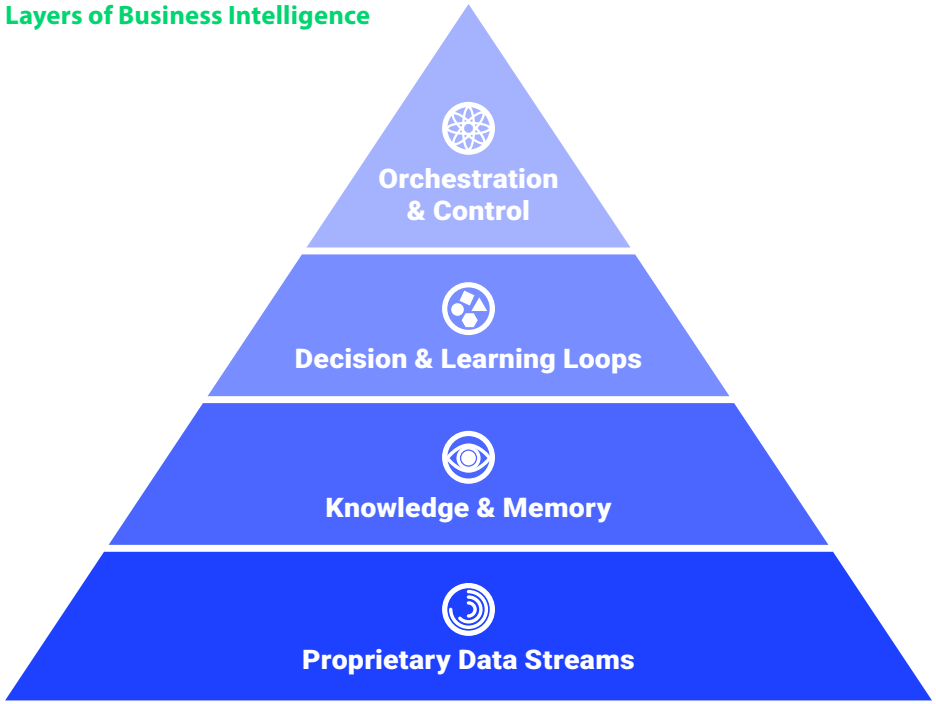

EPILOGUE

WHAT'S NEXT FOR YOUR BUSINESS?

Trying to chart a sensible AI roadmap amid today's breakneck technological evolution can feel a bit like dancing in a Ferrari — exhilarating yet disorienting. It's crucial to stay focused on what actually creates long-term competitive advantage for your business.

Lay the foundation with powerful data curation, paired with whole-brain fine-tuning. Rethink workflows from a blank sheet. Design optimal human–AI teaming that creates multiplier effects and unlocks value through user networks. These fundamentals make a world of difference over time.

If your organisation has launched AI pilots that aren't scaling or you're preparing your first enterprise-wide AI initiative, **the difference between a 5% and 95% successful outcome lies in avoiding these seven sins.**

THE ENTERMIND EDGE

7 IMPERATIVES TO WIN

With our collective experience driving 700+ enterprise AI implementations across Asia-Pacific, Latin America, Middle East and North Africa, Entermind is AI-native from day one, built to transform organisations from the ground up.

1. AI transformation is 90% orchestration, 10% technology
2. Every operational workflow deviation is a strategic advantage
3. Human-AI teaming must be redesigned from first principles
4. Federated AI governance outperforms centralised control
5. Business impact must be measurable from day one
6. We don't just deploy AI. We transform how your organisation works, learns and competes.
7. AI is the first whole brain technology in the world. It should be orchestrated using a whole brain approach that compounds engineering with empathy.

REFERENCES

1. Bellan, Rebecca. “Sam Altman says ChatGPT has hit 800M weekly active users.” TechCrunch, 6 Oct. 2025. Available at: https://techcrunch.com/2025/10/06/sam-altman-says-chatgpt-has-hit-800m-weekly-active-users/

2. MIT NANDA, State of AI in Business 2025: The GenAI Divide, July 2025. Available at: https://www.slideshare.net/slideshow/state-of-ai-in-business-2025-mit-nanda/282804851

3. AI4SP. Why 90% of AI Startups Fail. September 2024. Available at: https://ai4sp.org/why-90-of-ai-startups-fail

4. Wehbe, R. R., Lank, E., & Nacke, L. E. (2017). Left Them 4 Dead: Perception of Humans versus Non-Player Character Teammates in Cooperative Gameplay. In Proceedings of the DIS 2017 Conference (pp. 403-411). University of Waterloo, University of Lille, University of Ontario Institute of Technology. Available at: https://rinawehbe.ca/wp-content/uploads/p403-wehbe-1.pdf

5. SAP News Center. “AI in the Workplace Survey” August 27, 2025. Available at: https://news.sap.com/2025/08/new-walkme-survey-shadow-ai-rampant-training-gaps-undermine-roi

6. Bergdahl, J., Latikka, R., Celuch, M., Savolainen, I., Soares Mantere, E., Savela, N., & Oksanen, A. (2023). Self-determination and attitudes toward artificial intelligence: Cross-national and longitudinal perspectives. Telematics & Informatics, 82, Article 102013. Available at: https://doi.org/10.1016/j.tele.2023.102013

7. IBM. Calvesbert, E., & Raymond, B. (2025) “Unstructured: Turning unstructured data into AI-ready fuel”. Available at: https://www.ibm.com/new/announcements/ibm-and-unstructured-io-partner-to-accelerate-ai-ready-data-in-watsonx-data IBM

8. Borkan, A., AlYaychli, N. F., Mohammed, M., Safar, F., Omar, F. S., & Rasheed, M. H. (2025). Automating RLHF-based Hallucination Tracking. Journal Européen des Systèmes Automatisés, 58(7). https://doi.org/10.18280/jesa.580715

Inputs
Prashant Kumar
Sourabh Agrawal
John Woo
Jen Thong

With strategic input from our advisors and domain specialists across data, design, technology and organisational transformation. Thank you to the people who have made this whitepaper possible.

About Entermind
As the world’s first whole brain data and AI consultancy, we bring together a native understanding of the human experience with expert data practice across engineering, architecture and strategy to elevate business performance.

Make AI real with us

Write to us at
prashant@entermind.com

No AI was used to write this whitepaper.
Entirely written by humans.

7 Original Sins of AI Investment

A whitepaper by

Entermind