



Winning with AI



The business leader's guide to AI from
strategy to execution – an IMD playbook





Contents

Executive summary	04
The future of Generative AI: Emerging technologies and paradigm shifts	08
Understanding	10
GenAI marks a strategic inflection point – yet most leaders aren't ready	12
How GenAI implementation varies across functions	16
Productivity gains at work	18
Spotlight: Deploying GenAI to drive sustainability	22
Lessons from organizations that have effectively implemented AI: Defining AI maturity across financial services, consumer goods, automotive, healthcare and pharmaceuticals, and manufacturing	28
Building	32
The real AI challenge: Empowering people, not just platforms	34
The Value-Data-People framework: Applying IMD's model to real-world scenarios	40
Building for success: Capabilities and risk management	44
Managing risk	50
Identifying and mitigating potential pitfalls	52
Cybersecurity: How to build a secure-by-design strategy	54
Conclusion	58
Takeaways and recommendations	60

Executive summary

Understanding the AI landscape: Current trends and challenges in AI adoption

How much will the application of artificial intelligence (AI) change your organization over the next 12 or 36 months? According to [Gartner's 2025 predictions](#), by next year, 20% of organizations will use AI to flatten their structure, eliminating more than half of current middle-management positions, and within four years, one in 10 global boards will use AI guidance to challenge executive decisions that are of central importance to their business. Look at your company today: who you are, and how you make decisions, will shift dramatically in the next few years.

The adoption of generative AI (GenAI), in particular, represents a pivotal moment, with the potential to transform business as much as the advent of the internet. But along with this potential to reimagine how we work, create, and solve problems comes a significant responsibility for the humans using these tools to wield them in a way that benefits society. As the technology matures, some existing problems, such as data bias, will likely diminish. But others will arise. For example, the rapid adoption of AI comes at an environmental cost: a report from the US Energy Department last year found that AI may increase the portion of the nation's electricity supply consumed by data centers from 4.4% to 12% by 2028 in the US alone. To meet this demand, the report suggested, some power plants are expected to burn more coal and natural gas.

Value creation across industries

At IMD, we have seen a huge upswing in the operationalization of AI designed for autonomy; agentic AI in business goes beyond automating tasks, as we have seen in industrial AI for more than a decade; now it is building systems that manage and optimize whole business functions on their own.

In healthcare, AI is used in diagnostics, personalized treatment, drug discovery, and development. AI is deployed across financial services to reduce fraud, enhance security, and customize financial advice and services. If autonomous vehicles aren't yet the norm in logistics, they're on the rise, and smart cities everywhere rely on intelligent traffic management and congestion control. In IMD's annual [AIMaturity Index](#), we rank 200 of the world's leading companies

on how effectively they have adopted AI to transform their business strategies and operations. The results are [surprising](#). Yes, tech companies perform well. But representatives from financial services, telecoms, consumer goods, and the energy sector all make appearances near the top.

These organizations are implementing AI solutions, each in their own way and for their own reasons, but most importantly, they are all making progress. If you are a business leader, you have probably come to accept that AI is going to revolutionize many aspects of your work and personal life. However, you may be uncertain how to build a strategy to harness AI to its full advantage. You may be nervous about being left behind. Combined with our AI Maturity Index, this white paper aims to demonstrate what good looks like and provide a roadmap for effective deployment of AI.

Making the best investment

Many are optimistic about the ROI of AI. For every \$1 companies invest in AI, they are realizing an average of \$3.5 in return, according to a recent study conducted by IDC, and sponsored by Microsoft, with financial services, media, and telco having the highest demonstrated ROI... so far. The study found that one in 20 organizations worldwide realized an impressive \$8 in return. [Morgan Stanley Research estimates](#) that AI-driven productivity could add 30 basis points to 2025 net margins for members of the S&P 500. Others are far more conservative in their predictions: a [mid-2025 survey by Boston Consulting Group](#) found that for AI investments in the finance function, median ROI hovered at just 10%, with one-third of leaders surveyed reporting limited or no gains.

And yet, while ROI may vary, it is clear that firms are racing to invest: Gartner has predicted that worldwide GenAI spending could reach [\\$64.4bn in 2025](#), a massive increase of 76.4% from 2024. And timelines are shrinking: two in five organizations have implementation times of less than six months, according to the IDC study. These are high-stakes investments at speed.

This makes it critical that leaders are well-informed, spending in the right way, and mindful of the impact their decisions will have on their organizations, workers, and the planet. AI consumes a large amount of energy. There are also growing [concerns that AI](#)

has created a new digital divide between nations (mainly in the northern hemisphere) that have computer power and those that do not yet have the data centers to support AI. This divide will play a role in geopolitics, economic growth, and the ability to innovate. Now is the time to define the role that AI will play in your organization, prioritizing strategic objectives and clearly delineating the role that human agents will play in delivery.

The strategic imperative: Why AI transformation can't wait

Traditional AI focuses on analyzing data to make better predictions and improve decision-making. From B2C applications like Netflix's personalized content recommendations to industrial environments like Siemens' machine-learning-based predictive maintenance solutions, its use has altered the way we do business for more than a decade. We've also seen many companies continually struggling to extract value from AI for a variety of organizational, data, or capability reasons.

But we are now at another pivotal moment: GenAI has changed the game, enabling machines to create new content, such as text, images, music, and code, by learning patterns from existing data. GenAI is shifting the landscape of digital transformation by extending the capabilities of AI from analytical and decision-making tasks to creative and generative ones. This unleashes new possibilities for automation, personalization, and innovation within organizations. This playbook draws on the rich research from across the IMD faculty to provide an actionable guide to what leaders should be doing today to use the best of AI capabilities in their industries. From the C-suite to operational leaders across organizations, this paper explores the business case for, and addresses the risks of, embedding AI in your strategy.





The future of Generative AI: Emerging technologies and paradigm shifts

As organizations integrate GenAI, the technology landscape is rapidly evolving. Understanding these emerging trends is crucial for businesses preparing for the next wave of AI-driven transformation.

Hardware evolution: Accelerating AI performance

The computational demands of generative AI are driving innovation in specialized processors. Beyond traditional GPUs, dedicated AI chips like TPUs and NPUs are becoming more sophisticated, while neuromorphic computing mimics brain architecture for dramatically improved energy efficiency. Quantum-AI hybrid systems promise exponential speed improvements for certain problems. These advances will enable organizations to run complex AI models locally, reducing response speeds and improving privacy while making AI more accessible to smaller businesses.

Revolutionary model architectures and capabilities

Generative AI is expanding beyond text to multimodal systems that seamlessly process and generate text, images, video, and audio. Video generation models are approaching photorealistic quality, enabling organizations to create custom content at unprecedented scale.

New architectures like Joint Embedding Predictive Architecture (JEPA) represent a departure from traditional transformer models. JEPA enables more efficient learning through prediction rather than reconstruction, reducing computational requirements while improving reasoning capabilities. Mixture-of-experts models dynamically activate components based on tasks, making large-scale AI more efficient for organizations with diverse use cases.

Retrieval augmented generation: The key to accurate AI

Retrieval Augmented Generation (RAG) combines generative capabilities with real-time access to relevant, organization-specific information. Rather than relying solely on training data, RAG systems query databases and documents to provide accurate, contextually relevant responses.

For organizations, RAG addresses critical challenges: hallucination, outdated information, and a lack of domain expertise. By connecting AI to proprietary databases and real-time feeds, companies ensure accurate information while maintaining competitive advantages. RAG also provides audit trails crucial for regulated industries.

The rise of agentic AI: From tools to autonomous collaborators

Perhaps the most transformative development in GenAI is the evolution toward agentic systems – AI that can independently plan, execute complex multi-step tasks, and interact with various tools and systems to achieve objectives. Unlike current AI tools that respond to prompts, agentic AI can break down complex goals into actionable steps, execute those steps across multiple platforms and applications, and adapt its approach based on outcomes.

Current developments in agentic AI include systems that can autonomously manage email communications, coordinate project timelines across team members, conduct market research by interfacing with multiple data sources, and even negotiate contracts within predefined parameters. These systems represent an evolution from the reactive AI frameworks described in our earlier discussion to proactive AI collaborators.

The most significant paradigm shift lies in the transition from human-initiated to AI-initiated interactions. Instead of waiting for human prompts, future AI agents will proactively reach out to users with insights, recommendations, and alerts based on continuous monitoring of relevant data streams and changing conditions. This represents a fundamental change in the human-AI relationship, moving from a command-and-response model to a collaborative partnership where AI takes initiative in driving conversations and actions.

AI-initiated interactions: The next frontier

Future AI will proactively engage humans rather than merely responding to requests. AI agents will continuously monitor business conditions and team performance to identify opportunities, reaching out with insights and recommendations before problems escalate.

This fundamental transformation can be understood through the evolution of AI interaction patterns across organizational levels:

Level	Human-initiated	Machine-initiated
Individual	User prompts AI for specific tasks, information, or assistance	AI observes individual behavior and proactively provides personalized coaching, development suggestions, and performance insights
Team	Team members request AI to contribute to discussions, fact-check, or provide analysis	AI monitors team dynamics and autonomously coordinates work, assigns tasks, and optimizes team performance
Organizational	Humans set organizational goals, and AI executes complex multi-step processes to achieve them	Multiple AI systems independently identify strategic opportunities and initiate cross-organizational initiatives

The future of AI lies in the right column - systems that move from reactive response tools to proactive partners that observe, analyze, and initiate conversations and actions across all organizational levels. This represents a fundamental shift from *human prompts* → *AI responds to AI* observes → *AI initiates* → *human collaborates*.

This shift requires new frameworks for AI governance and human-AI collaboration. Organizations must establish clear boundaries for AI initiatives while maintaining oversight and transparency.

Preparing for the AI-native future

Companies should develop robust data governance for RAG implementations, establish policies for AI-initiated actions, and train teams for collaborative work with proactive AI systems. The organizations that will thrive view AI not as replacement technology, but as sophisticated collaborators capable of augmenting human intelligence and taking initiative in driving success.

Keep this framing in mind as you explore the rest of this white paper.



GenAI marks a strategic inflection point – yet most leaders aren't ready

The route to scaled business adoption of general-purpose technologies has historically been a long one, due to the organizational and work adaptation required to create and capture value. We are still developing an understanding of the transformative capabilities of generative AI (GenAI), but it seems likely to have a shorter business adoption cycle – and if deployed well, promises the strategic advantages of flattening organizational structure and turbocharging data-based decision-making.

Explore the current landscape now

As a leader, keep in mind three fundamental changes that characterize the phase of digital transformation we are currently in:

I: Geeks to crowd: The democratizing power of AI

To execute digital transformations, organizations traditionally have had to build AI capabilities. These units – centers of excellence, or AI factories – required skilled workers with deep expertise in data and computer science; these workers were rare, and therefore expensive. The units traditionally operated behind the scenes, organizing data, designing algorithms, and testing AI models, influencing outcomes with little direct user engagement until the system was ready for production. These experts sit at the top of the workforce pyramid.

With GenAI, the pyramid is inverted. It engages directly and interactively with users at all levels of the organization because there is no coding required, no algorithmic design, just multimodal and intuitive natural language interfaces. GenAI creates the potential for a large-scale, bottom-up, new wave of digital innovation. This presents both opportunity and risk. The opportunity lies in the wide accessibility of the technology, creating myriad potential applications for business value. But herein lies the risk, as applications can mushroom in every corner of the organization. A multitude of uncoordinated digital initiatives rarely generates a positive business return.

Action: Manage your GenAI initiatives within your existing digital transformation program; don't create a separate organizational governance if you have one that works. Focus on the few lighthouse applications that both advance your digital strategy and can move the (measured) needle on your business results. For the rest, "let the children play" within clear guardrails, and capture best practices.

II: Automation to augmentation: Building the right goals, operations, and skills

From substituting routine and repetitive tasks such as picking and packing in warehouses to automating banking back-office processes through software-based technology like robotic process automation, traditional AI is associated with automation, adopted to drive cost savings and improve efficiency. GenAI also automates tasks, but it pushes the frontier of automation to cognitive skills like analytical and creative activities; it can form hypotheses or generate product prototype ideas.

This has at least three profound effects on work roles. First, it allows employees to redistribute the time spent on certain tasks, focusing on those that add the most value to the organization. Secondly, it releases employees completely from some tasks to perform new ones. Thirdly, academic research in several industries has shown the upskilling and levelling effect of GenAI applications (i.e., lower-skilled workers benefit more than higher-skilled workers), allowing less skilled employees to perform tasks of higher complexity, which would normally be undertaken by more senior workers. This is where the new "future of work" frontier lies: the promise of higher worker productivity, efficiency, and quality output at scale.

Action: If the key strategic driver for your GenAI deployment is labor cost reduction, you will miss most of the potential for value creation. Instead, amplifying and augmenting workforce tasks and roles should be your north star. To realize the productivity benefits, work augmentation will demand process and workflow

redesign, and sometimes even business model evolution. Scaled GenAI implementation will change the way people and organizations work.

Human skills and competencies still matter in a GenAI world. Dynamic reskilling, experiential learning, and job evolution programs are critical to ensuring adoption and extracting full value from GenAI deployment. AI-augmented employees will drive a new human/machine division of labor within organizations, but it will have to be managed from a people-centric perspective to succeed.

III: Modernization to transformation: Adopting a culture of constant adaptation

Differentiation in the age of GenAI is not solely based on the technology, but more significantly on organizational execution. Start with ensuring superior data availability and accessibility (including proprietary), and the ability to derive analytical insights; both are critical for differentiation.

In traditional digital transformation, the focus has often been on adoption; with GenAI, the focus will be on constant adaptation. GenAI model lifecycles are short, and system capability improves faster than previous innovative digital technologies. Organizations that develop an edge for fast "sandboxing" of the various generations of the technology will speed ahead. Use-case design, data integration, testing, security, etc., will need to be streamlined and will not cope with the cross-silo approach of many traditional firms.

Finally, although GenAI technologies are widely accessible, the cost of developing the core models can be high. For non-tech organizations, it may be out of reach. And transformative use cases are, very often, not based on a single technology but on the integration of various systems.

Action: Organizations that are able to build and sustainably manage partner ecosystems to support their AI ambitions will get an advantage. Data, innovation processes, and ecosystems were always part of traditional digital transformation. What GenAI adds is the need for sustained and simultaneous excellence in integration, adaptability, and fast execution. Not a small challenge for digital leaders. The GenAI world offers low barriers to entry for modernization, but poses high execution hurdles

for true transformation. As in previous digital transformation phases, some organizations will graduate to the transformation stage, but many will be left behind.

For digital transformation, GenAI brings enormous potential for a step change in efficiency, productivity, and how work is performed. Understanding where the potential lies, and how to extract the greatest value for organizations, workers, and society, is top of the agenda for leaders right now. In the following sections, we will offer a roadmap, exploring what is possible, how to embed AI in your strategy across your organization, and how to build a culture of adaptation as technology develops.

Executive checklist

- Do we have the right people and learning programs for successful GenAI adoption across our organization?
- Do we have the organizational change management capabilities to redesign processes and workflows around human-AI collaboration?
- Are we managing GenAI initiatives within our existing digital transformation?
- Can we rapidly "sandbox" and adapt to new GenAI model generations faster than our competitors?



For digital transformation, GenAI brings enormous potential for a step change in efficiency, productivity, and how work is performed. Understanding where the potential lies, and how to extract the greatest value for organizations, workers, and society, is top of the agenda for leaders right now.

Didier Bonnet
IMD Professor of Strategy and Digital Transformation



How GenAI implementation varies across functions

To maximize the potential of generative AI (GenAI), organizations must think beyond individual uses – for example, drafting emails or managing projects – and explore how these tools can integrate into their broader strategy and across whole functions. Whether it's automating routine processes, fostering cross-team creativity, or creating deeply personalized customer experiences, the possibilities are transformative when GenAI is embedded across an entire organization.

It is not just large corporations that are benefiting from applying GenAI. Small- and medium-sized businesses can utilize GenAI to take on larger challenges and grow rapidly.

Driving efficiency across teams

Organizations often face the challenge of repetitive, time-consuming tasks that span multiple teams or departments. GenAI tools can help automate these tasks, for example by generating reports, summarizing meetings, or drafting routine communications. By deploying the capabilities at an organizational level, companies can significantly reduce the administrative burden on employees, freeing up time for more strategic activities. For example, a customer service team could benefit from AI-generated responses to common customer inquiries, while HR departments might use GenAI to streamline the onboarding process by automating documentation and training materials. These efficiencies across teams result in faster processes, reduced operational costs, and more cohesive workflows.

Enhancing innovation and product development

Beyond routine tasks, GenAI tools can play a key role in driving organizational innovation. By analyzing market trends, user behavior, and competitor products, GenAI can produce fresh ideas for product development, new features, and branding campaigns. Instead of relying solely on individual creativity, companies can apply these tools at scale, encouraging collective ideation sessions and cross-departmental brainstorming. Imagine a scenario where R&D, marketing, and design teams collaboratively use GenAI tools to generate prototypes, predict market response, and iterate ideas more quickly. This integration creates a unified approach to innovation, making it more efficient and less reliant on isolated inspiration.

Scaling personalization at the organizational level

Another key differentiator for organizational use of GenAI tools is the ability to deliver large-scale personalization. Marketing departments can deploy AI-driven systems to create unique content for thousands of users simultaneously, tailoring messaging for different customer segments based on their preferences and behavior patterns. Similarly, sales teams can utilize GenAI insights to deliver highly personalized pitches to potential clients, enhancing the overall customer experience. This kind of personalization, when applied at scale, helps organizations build deeper relationships with their customers and differentiate themselves from competitors.

Improving decision-making with data-driven insights

By analyzing vast datasets, GenAI tools can produce detailed insights, identify patterns, and create predictions that inform strategic choices along the entire decision-making pipeline. For example, a finance department could employ AI models to generate forecasts and assess financial risk, while operations teams could use these insights to optimize supply chain logistics. GenAI tools become invaluable when their predictive capabilities are aligned with organizational objectives, turning raw data into actionable strategies.

Solving emerging problems with scenario planning

Organizations are beginning to be able to process and respond to vast streams of data instantaneously, revolutionizing customer service, supply chain management, and operational decision-making. When we are confronted with new problems, the tool to use in the supply chain is very often Microsoft Excel. Today, GenAI is a new solution to new problems that appear in very volatile contexts. GenAI excels at scenario analysis and problem-solving. It allows supply chain managers to model different outcomes and quickly adapt strategies based on shifting conditions. This flexibility is especially important in a world where disruptions are increasingly the norm. After years of experimentation and frustration, supply chain professionals are finally seeing the results they hoped for from GenAI. Tools that once seemed like science fiction are becoming practical and accessible, offering new ways to optimize operations, reduce costs, and navigate uncertainty.

Filling the talent gap

Supply chain management and other functions often struggle to attract skilled professionals, and many tasks remain manual and repetitive. GenAI can automate these tedious processes, freeing up teams to focus on higher-value, strategic work. Looking ahead, the development of AI agents capable of making autonomous decisions could further streamline operations. These agents promise to handle routine tasks independently, reducing the burden on human operators and increasing efficiency.

Strengthening knowledge sharing and collaboration

GenAI can serve as a bridge between siloed departments by summarizing, cataloging, and making organizational knowledge more accessible. AI-driven platforms can be used to transform even unstructured data, such as meeting notes, research findings, and departmental documentation, into easily digestible formats by helping teams across the company stay informed and aligned. Unlike individual use, where an employee might generate a single document or summary, deploying GenAI at scale ensures that all organizational knowledge is structured, stored, and accessible in a consistent manner – breaking down barriers to effective collaboration.

When used at scale, GenAI tools have the power to fundamentally reshape workflows, boost innovation, and create personalized experiences that were previously unachievable. By adopting these tools strategically, organizations can unlock efficiencies across departments, drive coordinated innovation, and make smarter, data-informed decisions – positioning themselves for sustained success in an increasingly competitive market.

Productivity gains at work

Personal gains

Application area	Key points
Productivity	Automates repetitive tasks by creating a custom agent to verify fields and draft follow-up emails. Automates tasks like meeting summaries, highlighting key themes, and suggesting follow-up actions.
Communication	Drafts personalized messages by analyzing context, tone, and prior exchanges.
Creativity	Brainstorming, generating novel ideas aligned with discussions. Helps teams overcome creative blocks, speeding up ideation and planning processes.
Project management	Identifies bottlenecks and potential risks from team conversations. Proactively suggests solutions, adjusts timelines, and monitors project dynamics.
Strategic decision-making	Processes vast information for context-aware recommendations. Adapts suggestions based on real-time feedback, acting as a virtual advisor.

Organizational gains

Gain	How
Driving efficiency across teams	Automates repetitive tasks like report generation, meeting summaries, and onboarding. Benefits include faster processes, reduced costs, and streamlined workflows.
Enhancing innovation	Facilitates ideation by analyzing market trends, user behavior, and competitor products. Encourages cross-team brainstorming and faster prototyping for efficient innovation.
Scaling personalization	Enables large-scale personalized marketing and sales strategies. Tailors customer experiences across segments, building deeper relationships and competitive differentiation.
Data-driven decision-making	Analyzes large datasets for insights, patterns, and predictions. Supports strategic decisions in finance, operations, and overall organizational planning.
Strengthening knowledge sharing	Summarizes and catalogs organizational knowledge – breaking down silos, making unstructured data accessible and improving collaboration and alignment across departments.
Scaling GenAI impact	Transforms workflows and boosts innovation across entire organizations, drives coordinated, efficient, and personalized strategies for competitive advantage.

Executive checklist

- Am I deploying GenAI to drive personal productivity and creativity?
- Does our GenAI strategy extend beyond individual efficiencies to deploy capabilities across entire functions – and the whole organization?
- Are we enhancing collaboration across departments and functions by ensuring that organizational knowledge is structured, stored, and accessible in a consistent manner?
- Are we using GenAI to drive innovation systemically?



Whether it's automating routine processes, fostering cross-team creativity, or creating deeply personalized customer experiences, the possibilities are transformative when GenAI is embedded across an entire organization.

Michael Wade
TONOMUS Professor of Strategy and Digital at IMD

Amit Joshi
IMD Professor of AI, Analytics and Marketing Strategy



Spotlight: Deploying GenAI to drive sustainability

If AI is widely recognized as a driver of productivity and efficiency, sustainability is often seen as a cost burden. However, by integrating AI into sustainability initiatives, organizations can optimize operations and accelerate the adoption of low-carbon technologies. This combination not only improves efficiency but also fosters the development of new business models that align environmental responsibility with economic growth. The urgency for this approach is growing, fueled by regulatory changes, investor expectations, and resource scarcity.

Generative AI (GenAI) is revolutionizing how businesses tackle their most pressing sustainability challenges. However, this transformative power comes with a significant cost: GenAI's immense energy and water demands pose serious environmental concerns. Despite these challenges, there are four areas where GenAI is not just improving processes, but fundamentally redefining how businesses can and should lead in a more sustainable world.

I. Task: Improve sustainability reporting

Sustainability reporting is a strategic imperative for businesses, but frameworks like CSRD, CDP, ISSB, GHG Protocol, TCFD, GRI, and others demand substantial effort and precision. For many organizations, the manual processes traditionally used for data collection and analysis are no longer manageable.

Action: By deploying GenAI to manage that complexity, you can automate data collection across your entire supply chain, synthesizing information from supplier disclosures, IoT sensors, and satellite imagery into a unified, actionable dataset.

Advanced AI-powered solutions, like those developed by Accenture and Avanade, take this a step further by seamlessly integrating qualitative and quantitative data. Quantitative metrics are automatically imported and aligned with regulatory standards, while GenAI

assists in drafting qualitative responses to ensure narratives are both compliant and forward-looking. These systems also adapt dynamically to multiple frameworks, allowing organizations to manage diverse reporting obligations without duplication of effort.

With integrated dashboards offering real-time insights, companies can visualize trends, monitor progress, and uncover performance gaps. By automating complexity, GenAI transforms sustainability reporting into a streamlined, value-driven process, enabling leaders to focus on strategy and impact rather than administration.

Impact: Tasks that once took weeks or months can now be completed in hours. GenAI for sustainability reporting also achieves unparalleled accuracy, so organizations can identify risks and inefficiencies early, turning reporting from a compliance exercise into a proactive tool for decision-making.

II. Task: Optimize resource use

GenAI is revolutionizing how organizations manage their most critical inputs, enabling leaders to achieve more with less while uncovering new pathways to sustainability.

Action: Deploy GenAI to react to multiple data inputs to drive smarter use of resources. For example, in agriculture, by analyzing data from sensors, drones, and satellite imagery, GenAI tailors interventions to the specific needs of each plant, whether it's water, nutrients, or pesticides.

In manufacturing, GenAI monitors every step of production, predicting defects before they occur, and automatically adjusting machine settings to minimize material waste and energy use. Companies like Siemens are using GenAI-driven systems to refine production processes in real time, achieving unprecedented levels of efficiency and sustainability.



Impact: Building scalable, hyper-localized approaches eliminates the inefficiencies of blanket treatments, conserving resources while improving outcomes. This isn't just about fine-tuning; it's about the transformation of static systems into adaptive ecosystems, where decisions are data-driven, operations are seamless, and every resource is used to its fullest potential.

III. Task: Drive circularity

From revolutionizing recycling to designing for reuse and optimizing waste conversion, it's shifting the narrative from waste reduction to value creation. Future-oriented leaders know that transitioning to a circular economy, where resources are reduced, reused, refurbished, or recycled, isn't just the right thing to do; it's a strategic necessity. GenAI is accelerating this shift by enabling businesses to reimagine waste not as a liability, but as a valuable asset.

Action: At the beginning of the life cycle, GenAI is redefining product design by embedding circularity into the very DNA of new products. AI systems analyze material properties, consumer usage patterns, and end-of-life scenarios to optimize products for reuse and recycling. Think modular furniture that can be disassembled and refurbished, or consumer electronics like the [Fairphone](#), which is engineered for easy upgrades and repair instead of replacement.

In addition to product design interventions, GenAI is transforming waste-to-energy systems, optimizing processes to extract the maximum energy output while minimizing emissions. These systems don't just mitigate waste, they convert it into clean, usable energy, closing the loop on resource use and turning environmental responsibility into operational efficiency.

Impact: At the end of the life cycle, AI-powered sorting technologies use advanced computer vision to identify materials, even in mixed waste streams, with unparalleled accuracy. These systems dramatically reduce contamination, ensuring higher-quality recyclables that fetch better prices in secondary markets. More importantly, they minimize reliance on virgin materials, reducing extraction costs and environmental impact. What was once "waste" is now a revenue stream.

IV. Conserve biodiversity

GenAI is turning the tide of biodiversity conservation. It provides the tools to monitor, protect, and restore ecosystems at a scale and precision never before possible. For businesses, this is a chance to go beyond damage control and lead the way in safeguarding the planet's natural capital.

Action: Use GenAI to drive conservation, empowering businesses to protect and sustain the natural systems they depend on. From AI-powered drones surveying vast rainforests, analyzing live footage to detect illegal logging, poaching, or habitat loss. These systems, already deployed in regions like the Amazon, provide real-time insights into threats that might otherwise go unnoticed. Tools like OceanMind monitor global fishing activities, using satellite data and AI to identify illegal practices and prevent overfishing.

In addition to critical monitoring activities, GenAI enables restoration, analyzing ecological data to predict the success of reforestation projects and habitat restoration efforts.

Impact: Conservation efforts that once relied on reactive measures are now proactive, allowing businesses and governments to intervene before damage is done. By ensuring the integrity of critical ecosystems, companies protect not only biodiversity but also the long-term viability of their own supply chains.

Executive checklist

Does my organization integrate AI into its sustainability strategy across:

- Disclosure and non-financial reporting?
- Resource optimization?
- Design and research & development?
- Safeguarding natural capital?

Using GenAI to build a more sustainable business

Task	Key points
Sustainability reporting	Analyzes data, automating across all sources, to extract strategic insights.
Optimizing resource use	Enhances efficiency, transforming static systems into dynamic and adaptive ecosystems.
Designing for circularity	Drives circularity, revolutionizing recycling and designing for reuse, and optimizing waste conversion.
Conserving biodiversity	Analyzes vast datasets to allow prompt proactive intervention before the damage is done.



Tasks that once took weeks or months can now be completed in hours. GenAI for sustainability reporting also achieves unparalleled accuracy, so organizations can identify risks and inefficiencies early, turning reporting from a compliance exercise into a proactive tool for decision-making.

Julia Binder
IMD Professor of Sustainable Innovation and Business Transformation

Lessons from organizations that have effectively implemented AI: Defining AI maturity across financial services, consumer goods, automotive, healthcare and pharmaceuticals, and manufacturing

IMD's **TONOMUS Global Center for Digital and AI Transformation** provides world-leading [research](#) and analysis on digital transformation and prepares executives to take advantage of digital opportunities, especially AI. These insights also inform IMD's digital and AI transformation programs for individuals and organizations.

As part of its annual AI Maturity Index, the center has created this checklist to evaluate how effectively organizations are leveraging AI technologies to transform their business operations and strategies:

1. Strong executive support

AI maturity starts with strong leadership commitment and executive involvement. This is demonstrated by a declared corporate AI strategy, the appointment of AI-related C-suite roles, and the establishment of a dedicated AI business unit. Such strategic initiatives demonstrate a company's dedication to integrating AI into its core business operations.

2. Technology and infrastructure

Investing in scalable, cloud-based data platforms is crucial for AI maturity. These platforms facilitate seamless data access and manipulation, ensuring AI systems are integrated across business units. This integration supports operational workflows and drives innovation in AI-powered products and services, creating new value propositions and enhancing customer experiences.

3. Operational excellence


AI maturity involves leveraging AI for automation and integrating it into organizational processes to enhance decision-making and improve productivity. Organizations demonstrate the practical application of AI in achieving operational excellence, showcasing declared uses and examples of AI in their processes.

4. Workforce development and culture

Continuous investment in reskilling and upskilling employees ensures that the workforce can work effectively with AI technologies. This fosters a culture that embraces AI. Investments in AI skills development and strategic hiring practices are key measures of progress in this dimension.

5. Ethics and risk management

Achieving AI maturity necessitates the responsible use of AI, which hinges on robust ethical governance. To this end, organizations establish ethics and governance committees, develop guiding principles, and create structures for responsible AI use. Additionally, they launch and sustain initiatives focused on responsible AI practices, ensuring ethical considerations are embedded in their AI operations.



AI by industry



Financial services

The most successful institutions are those that have systematically developed their AI capabilities across multiple dimensions, creating sustainable competitive advantages through technological innovation paired with organizational transformation.

➔ Read more about industry leaders, including Visa (#12 in the AI Maturity Index), Barclays (#21), and JPMorgan Chase (#29) here.



Automotive

Artificial intelligence is reshaping every aspect of the automotive industry, from manufacturing operations to autonomous vehicle deployment. Industry leaders recognize that AI maturity extends beyond isolated use cases to encompass comprehensive organizational change, transforming how companies design, build, and sell vehicles.

➔ Read more about industry leaders, including Volkswagen Group (#22 in our AI Maturity Index), Hyundai Motor (#23), Toyota Motor (#25), and General Motors (#27), here.



Manufacturing

Manufacturing has entered a new era where artificial intelligence is becoming the deciding factor determining competitive advantage. Industry leaders understand that AI maturity involves comprehensive organizational transformation, revolutionizing everything from production processes to customer engagement across global operations.

➔ Read more about industry leaders, including Lockheed Martin (#30 in the AI Maturity Index), General Electric (#31), Contemporary Amperex Technology (#35), Mitsui & Co. (#40), and Siemens (#41), here.



Consumer goods

Artificial intelligence has become a defining factor in the consumer goods sector, with forward-thinking organizations systematically developing capabilities that transcend individual projects to create enterprise-wide transformation. These companies recognize that AI maturity isn't just about implementing technology – it's about reimagining their entire business model.


➔ Read more about industry leaders, including Sony (#18 in the AI Maturity Index), Walmart (#24), and The Home Depot (#75), here.



Healthcare and pharmaceuticals

Organizations at the forefront of this transformation are building AI capabilities across multiple areas of operation, establishing market advantages through technological innovation and organizational transformation. Healthcare organizations have established specialized technical infrastructures that drive innovations in their products, research capabilities, and therapeutic technologies and have embedded AI into their operational processes and workflows, focusing on applications that improve efficiency, reduce costs, and enhance business performance. Industry leaders extend their capabilities through specialized partnerships that accelerate development while addressing medicine's unique collaborative research tradition.

➔ Read more about industry leaders, including Bayer (#20 in the AI Maturity Index), Medtronic (#26), and AstraZeneca (#28), here.



Executive checklist

- Does my organization's strategy integrate AI into core business operations?
- Does my organization have quantified productivity targets for our AI initiatives?
- Have I set my workforce up to embrace AI through regular upskilling and reskilling programs?
- Does my organization have clear ethics and governance guidelines to enable responsible AI use?

Building

The implementation
blueprint: From strategy
to execution

03



By Howard Yu and Mark Greeven



The real AI challenge: Empowering people, not just platforms

Now that we have explored the gains possible from AI and generative AI (GenAI) adoption, we will examine the best ways to build these technologies into your organization. AI's transformative potential requires more than technical readiness – it demands workforce readiness. Many organizations invest in digital learning platforms, but employees still struggle to apply new skills on the job. This isn't a training failure; it's a structural one.

The future belongs to organizations that ignite a sense of urgency around AI, while building the practical infrastructure to support experimentation and application. The new advantage lies in problem framing, judgment, and the ability to work symbiotically with machines – skills that emerge not through passive instruction, but through embedded, real-world experience.

Companies that fall behind in terms of skills will fall behind in terms of performance. They also risk irrelevance if their workforce isn't adapting to AI-enhanced work. To unlock that value, companies must rethink how people learn, practice, and adapt. Those that reskill successfully will capture the speed and scale advantages that companies like [Nvidia](#) now enjoy.

When learning falls flat

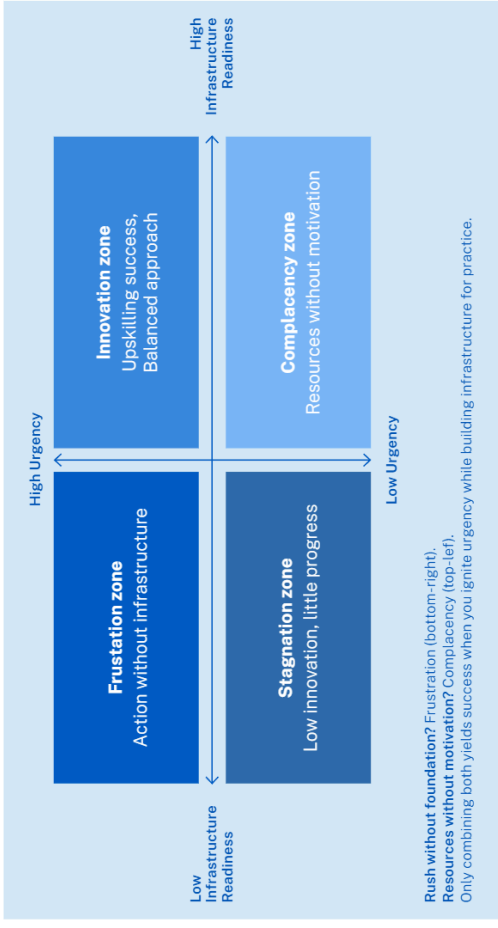
Companies are pouring resources into upskilling and reskilling, but many still struggle to turn training into tangible change. Employees often emerge from training sessions energized – brimming with ideas and eager to put new skills to work. But back at their desks, that momentum fizzles out. Without clear pathways to apply what they've learned in real situations, the training 'rush' fades fast. The result? A frustrating gap between potential and progress.

Reskilling isn't about ticking off e-learning modules or collecting certifications. It's about enabling people to drive change, whether by identifying opportunities to use AI or experimenting with solutions in their day-to-day roles. The jobs of the future will reward not rote execution, but critical thinking and iterative problem-solving in partnership with machines. The real transformation comes not from competing with AI, but from learning to thrive alongside it. Why do so many companies fall short? Because the barriers are cultural and structural. Hierarchies remain rigid. Siloed departments stifle experimentation.

Leadership hesitates to decentralize decision-making. In the absence of psychological safety, few employees will take the initiative to experiment with AI, regardless of training. The chart below illustrates a simple truth: reskilling only succeeds when infrastructure and motivation align. Pouring resources into training without empowering people to act leads to frustration. Cultivating urgency without support leads to chaos. True transformation comes when employees are both equipped and energized to apply what they've learned.

What drives successful reskilling

One chart reveals how rushed AI adoption backfires – and what actually works.



A modular design for flexibility

In the context of AI-driven transformation, modular design offers "innovation without the chaos." Take the case of [Mastercard](#), which shows how modular systems can drive digital innovation in the global payments industry. At its heart, modularity means breaking down complex, monolithic systems into modules, which enables faster modification and more efficient scaling.

Few areas of financial services have seen as much technological disruption as payments. The rise of digital and mobile tools has triggered the most fundamental transformation since the introduction of plastic cards half a century ago. Once dominated by a few banking behemoths, payments can now be made using a range of devices and across multiple channels.

Rather than relying on its reputation and affluent customer base, Mastercard responded by repositioning itself as a technology partner with a central commitment to modularity. This approach reshaped both Mastercard's internal systems and the way it engaged with partners.

That openness extended to the FinTech sector.

Mastercard has worked hard to position itself as a preferred collaborator for new entrants by sharing tools, offering access, and encouraging co-creation. Developers could plug into its network via application programming interfaces (APIs), underlining the strengths of a platform-based model. This model feeds a cycle of innovation that delivers new services to consumers and businesses alike.

From adoption to absorption

All companies can access frontier tech. But not all can turn it into productivity. Analysis by the OECD has found a growing gap between top-performing firms and the rest, even when access to frontier technologies is broadly the same.

The difference? Leading firms don't just adopt tech – they absorb it. They have complementary capabilities: advanced skills and workforce training, strong management and leadership, modern organizational practices, and access to finance, scale, and networks. These aren't nice-to-haves; they're what make the tech work.

Laggard firms often miss those pieces. They may have the same access to technology, but without the capacity to reshape workflows or upskill teams, the benefits stall. Same tech, different outcomes.

Another trap is the absence of robust metrics. By diligently measuring AI's real-world impact on cycle times, output quality, or customer satisfaction, leaders can ensure they stay ahead of the game. This means benchmarking AI-enabled workflows against traditional processes. Are customer service agents responding faster? Are supply chain forecasts more accurate? Is AI-enhanced personalization driving higher conversion rates?

At global learning company Pearson, CFO Sally Johnson isn't just talking about AI experiments; she's also communicating real impact. GenAI tools built into the company's higher education courseware are driving performance where it counts. "For products using those AI study tools, we've seen double-digit revenue growth," she says. "So, it's making a difference from a commercial point of view, as well."

A framework for AI adoption and implementation

To scale from isolated pilots to transformation, leaders must provide an organizational environment that welcomes AI.

Step 1: Invest in hands-on reskilling programs

Programs must go beyond passive learning. Employees should work on real projects that integrate AI tools, whether experimenting with predictive analytics or GenAI in content creation. The emphasis must be on "learning by doing" and business relevance.

Step 2: Create "sandboxes" in which employees can test AI innovations

Provide safe environments in which teams can test AI use cases without performance pressure. These sandboxes encourage experimentation and foster a culture of curiosity. Crucially, businesses must shield these sandboxes from bureaucratic drag to avoid smothering innovation.


Step 3: Establish metrics to track AI's internal and external impacts

Track how AI affects both internal efficiencies and customer-facing outcomes. Measuring impact helps maintain focus, justify investments, and reinforce a results-oriented mindset.

The flywheel effect

Effective reskilling creates a virtuous cycle, or the "flywheel effect." Learning feeds application, which leads to the results that motivate endeavor and accelerate investment.

Organizations that fail to act decisively will soon find themselves at a disadvantage in terms of talent retention, while struggling to maintain market relevance. As AI evolves, so does the competition. In a world racing to harness intelligent machines, the differentiator won't be access to technology; rather, it will be human adaptability. CEOs who reskill their people and empower them to act will set a pace that others will struggle to match.





In a world racing to harness intelligent machines, the differentiator won't be access to technology; rather, it will be human adaptability. CEOs who reskill their people and empower them to act will set a pace that others will struggle to match.

Howard Yu
IMD LEGO® Professor of Management and Innovation

Mark Greeven
IMD Professor of Management, Innovation



The Value-Data-People framework: Applying IMD's model to real-world scenarios

Ensuring AI adoption and absorption requires the above-discussed cultural shift; it also requires that C-suite leaders ensure that AI initiatives align with overarching business goals, avoiding divergent, siloed efforts, and driving competitive advantage.

Many AI projects fail to meet their objectives – often significantly underperforming. To navigate these complexities, organizations should focus on three essential dimensions of AI adoption: business value, data, and people. Together, these elements form a value-data-people framework, a structure for conceptualizing the critical questions that decision-makers must address. What value does the organization aim to create with AI? Does it have access to the required data? And how will employees and stakeholders perceive and adapt to the changes?

By keeping these considerations in focus, leaders can better prioritize resources, mitigate risks, and increase the likelihood of long-term success.

1. Value: Defining the business case

The first dimension of the framework challenges organizations to articulate the value they intend to create with AI. While this may seem obvious, many struggle to provide a straightforward answer when asked about the specific problem they are trying to solve.

Rather than adopting AI for its own sake, successful organizations focus on solving measurable challenges – such as using algorithms to improve sales performance. For example, a salesperson could use AI to predict which clients to approach, thereby increasing revenue from \$1m a year to \$1.3m – a specific use case that demonstrates tangible business value.

Companies that thrive with AI adoption tend to take a focused, pragmatic approach. They prioritize solving specific, manageable problems, accumulating small wins, and avoiding costly failures.

Incremental successes not only improve outcomes but also build internal momentum for more ambitious initiatives. The paradox here is that if leaders focus on funding initiatives aimed at solving tangible problems, they will indirectly be contributing to the generation of knowledge within the organization. This approach could even spark the cultural evolution needed to utilize AI to tackle larger, more complex problems further down the line.

2. Data: Ensuring access and quality

The second dimension of successful AI adoption focuses on data. Commentators often summarize this aspect using the mantra “garbage in, garbage out.” AI’s effectiveness depends on the quality and accessibility of data, yet organizations often lack the volume, diversity, or structure required for effective AI training. The key is to look at data in terms of access rather than just ownership. You may “own” data, but cannot use it because of a lack of consent. But there is also data you do not own but can access.

Data collaboration platforms enable organizations to train AI models while safeguarding privacy. These systems work by sending algorithms to where the data is stored, rather than moving data into the organization to train the tool. This ensures personal information remains securely stored at its source without impeding analysis.

Such platforms range from proprietary services offered by private companies to open-source solutions used by organizations or consortia. The widespread use of such platforms underlines the growing recognition of the value of securely tapping into shared or sensitive data. Importantly, these tools can address a critical challenge for AI development: the lack of high-quality training data. For example, hospitals and pharmaceutical companies can collectively train algorithms to support enhanced diagnostics or treatments without sharing raw data.

In more complex B2B environments, where regulations or privacy concerns prevent companies from using customer data to train AI models, these platforms allow firms to train algorithms while being impeccable in respecting privacy and facilitating compliance and innovation.

By maintaining privacy while enabling insights, data collaboration platforms are unlocking new possibilities across industries, from healthcare to autonomous vehicles, while navigating the growing complexities of data regulation. Without the relevant data, an organization’s focus should shift to acquiring or gaining access to it rather than just pressing ahead.

3. People: Managing perceptions and building trust

The third dimension – people – often determines whether an AI initiative succeeds or fails. AI can be seen as a threat, particularly due to concerns about job displacement. This reflects wider public anxiety. A Pew Research Center study found growing concern about the role of AI in public life, with 52% of US respondents saying they feel more concerned than excited about the increased use of AI.

Organizations must address these fears head-on, emphasizing that AI can enhance human capabilities and empower people to do more. The key message is that even with AI in place, human expertise will remain essential, especially as new challenges emerge. It is important to manage how employees perceive this transition – if they see AI as a threat, they may resist or even undermine initiatives.

Successful AI initiatives focus on communication and change management, recognizing that the wrong perception of AI can significantly increase the risk of failure. Executives must engage stakeholders early and often, gaining buy-in and ensuring alignment and trust throughout the transition process.

Executive checklist

- Have we identified the specific roles that will be most transformed?
- What is the value we want to create?
- Do we have access to the relevant data?
- How will our workforce perceive this change?
- Have we built a sandbox environment for safe testing?



Companies that thrive with AI adoption tend to take a focused, pragmatic approach. They prioritize solving specific, manageable problems, accumulating small wins, and avoiding costly failures.

José Parra Moyano
IMD Professor of Digital Strategy



Building for success: Capabilities and risk management

AI-based tools have already reshaped how businesses operate, yet many business leaders struggle to grasp AI's diverse abilities and how they might evolve in the years to come.

One impediment to success is the typical framing of AI progress, which is often explained through a series of levels or stages, when in fact, AI development does not move in a predictable sequence; real-world AI development is far more complex and dynamic.

To navigate this complexity, we need an AI capability framework that incorporates the potential for non-linear growth, refines what we mean by social and emotional intelligence, and clearly distinguishes advanced autonomy from the more speculative notion of AI consciousness.

Seven overlapping zones of AI capability: A leader's framework

Level 1. Foundational automation and pattern recognition

AI at this level excels at performing repetitive tasks and recognizing simple, predefined patterns. The primary focus of this foundational level is automating workflows with deterministic outcomes that can be easily modeled.

Business examples:

- **Process automation (RPA).** Tools that extract data from documents and input it into back-end systems, drastically reducing manual labor.
- **Basic chatbots.** Simple "FAQ bots" that follow scripts to respond to routine inquiries.

Level 2. Rule-based decision-making

This level leverages explicit logical rules to make decisions. Systems rely on predefined "if-then" conditions or structured logic trees, codifying human expertise into predictable frameworks. Rule-based decision-making tools excel in environments where the parameters are well-defined and ambiguity is minimal.

Business examples:

- **Tax preparation software.** Rule-based programs that guide users through a series of questions.
- **Workflow automation.** Enterprise tools that trigger approvals or escalate issues based on pre-set rules.

Level 3. Contextual adaptation and situational awareness

AI systems at this level use data-driven models, such as machine learning, to adapt to new or uncertain contexts in real time. Unlike rule-based systems, they can manage variability and anomalies not explicitly programmed into their logic. These systems excel at interpreting complex datasets and can "learn" from past interactions, making them highly valuable in industries requiring real-time adaptation. However, their proficiency is typically confined to a specific domain, and their ability to interpret social or emotional contexts remains limited. The interplay between levels becomes evident here, as these systems may demonstrate advanced behavior in isolated tasks while remaining rule-bound elsewhere.

Business examples:

- **Personalized recommendation engines.** Systems like those used by streaming platforms (Netflix) or e-commerce sites (Amazon) to adapt suggestions based on individual user data.
- **Dynamic supply chain management.** AI tools that automatically adjust shipping routes or inventory in response to demand fluctuations or real-time logistics data.

Level 4. Domain-specific social/emotional recognition

AI at this level can identify social and emotional cues within a defined context. Using modalities like text, voice, or facial analysis, these systems can infer user sentiments such as frustration, sadness, or satisfaction. However, their effectiveness relies on narrowly scoped applications, often tailored to specific industries or user groups. The capabilities of Level 4 systems are inherently domain-specific, meaning their success depends on carefully curated training data and clear boundaries. While they may appear empathetic in some contexts, their understanding of emotions is often shallow, rooted in pattern recognition rather than true contextual awareness. Misalignment between their perceived and actual competencies can erode trust if these tools are overextended beyond their intended scope.

Business examples:

- **Customer service chatbots with sentiment analysis.** When a caller expresses anger, the AI apologizes and escalates to a human if necessary.
- **Employee mood trackers.** These tools scan emails or Slack channels for signs of stress or dissatisfaction.

Level 5. Generalized social/emotional understanding

At this stage, AI systems begin to recognize and respond to social and emotional dynamics in diverse contexts. They can adapt their behavior based on individual backgrounds, cultural norms, and situational subtleties, offering nuanced and contextually appropriate interactions. Achieving generalized emotional intelligence is a formidable challenge. Success at this level requires bridging linguistic, cultural, and contextual gaps. Ethical concerns grow sharper, especially regarding data privacy and the potential misuse of emotionally intelligent AI for manipulative purposes.

Business examples (largely experimental):

- **Advanced counseling bots.** Systems that can differentiate between various forms of distress and modulate responses accordingly. Note that a recent Stanford study found these bots may not lack effectiveness when compared to human therapists, but may also introduce harmful stigma and dangerous responses.
- **High-level negotiation tools.** AI that handles multi-party negotiations with empathy and strategic insight, factoring in emotions, cultural norms, and stakeholder motivations.

Level 6. Advanced autonomy and strategic goal setting

This level represents a transition from task-focused AI to systems capable of independent goal setting and strategy formulation. These systems can operate with minimal human intervention, devising and executing plans that adapt dynamically to feedback or evolving circumstances. Advanced autonomy fundamentally shifts AI's role from operational support to strategic partner. These systems often function without direct oversight, raising critical questions about governance, accountability, and ethical responsibility. Interestingly, systems at Level 6 may operate effectively without any form of emotional intelligence, emphasizing that advanced cognition and emotional understanding are not inherently linked.

Business examples:

- **Autonomous trading systems.** AI that sets its own objectives ("maximize returns under these risk constraints") and dynamically trades without real-time human oversight.
- **Adaptive logistics.** A fleet of self-driving trucks that coordinates routes and loads, and schedules autonomously, responding to traffic or weather changes in real-time.

Level 7. Consciousness and self-awareness

This hypothetical level refers to AI systems with subjective experience or introspective awareness. Such systems would have a sense of self, enabling them to reflect on their goals, actions, and internal states in ways analogous to human consciousness. The prospect of AI consciousness is deeply speculative and often misinterpreted in popular discourse. While it may captivate imaginations, consciousness is not a prerequisite for transformative AI applications. Organizations should focus on the implications of advanced autonomy (Level 6), which presents more immediate challenges and opportunities. Misunderstanding the distinction between autonomy and consciousness risks ethical and operational missteps in AI deployment.

Business examples (purely speculative):

- None in current commercial or research settings is definitively conscious. It is an open question whether machines can ever achieve true self-awareness.
- A hypothetically conscious AI might argue for its rights, make moral judgments beyond its programming, or demonstrate unpredictable motives stemming from internal subjective states.

Non-linear pathways in action

A crucial part of this framework is recognizing that AI systems:

- Can skip levels. A large language model might show glimpses of domain-specific emotional recognition (Level 4) but still use basic rule-based or pattern-based logic (Levels 2–3) for other tasks.
- Can combine levels. You might see a somewhat autonomous system (Level 6) in a narrow application, yet it has only superficial emotional intelligence (Level 4).
- Do not proceed strictly upward. An AI's performance can plateau or even degrade when encountering unfamiliar environments.

Practical implications for business leaders

AI's rapid evolution through the seven levels outlined above carries profound implications for business strategy and leadership. Successfully navigating these changes requires continuous evaluation, risk management, ethical foresight, and proactive workforce planning.

To effectively leverage this framework, leaders should regularly assess how their AI systems perform across the seven levels, recognizing that AI tools often exhibit competencies that span multiple levels. Instead of assigning a single level to an AI system, it may be more insightful to create a capability matrix that maps the AI solution's performance across various dimensions. This approach helps uncover surprising strengths or weaknesses and guides targeted improvements. For example, conducting quarterly audits can track how well AI tools adapt to new data or changing conditions. These reviews might reveal areas where a system has unexpectedly advanced to a higher level or, conversely, regressed due to shifts in its operational domain.

As AI systems grow more autonomous (Level 6), the potential consequences of their errors also become more significant. For instance, a self-directed trading bot that misjudges market conditions could trigger significant financial losses almost instantly. Mitigating these risks requires robust governance frameworks and proactive monitoring.

To ensure accountability, organizations must document who is responsible for each AI system's design, deployment, and ongoing monitoring. Establishing intervention triggers is another critical step. These predefined thresholds – financial losses, safety incidents, or user complaints, for example – should automatically prompt human oversight or intervention. Additionally, compliance and governance efforts must stay ahead of regulatory developments, particularly as explainability becomes a key requirement for advanced AI. Leaders should ensure that they can clearly articulate how their systems arrive at decisions, fostering trust and meeting legal obligations.

Deploying AI systems with higher social and emotional intelligence (Levels 4–5) raises significant ethical considerations. These systems often interact directly with users in sensitive ways, necessitating transparency, fairness, and human oversight.



“ AI’s rapid evolution carries profound implications for business strategy and leadership. Successfully navigating these changes requires continuous evaluation, risk management, ethical foresight, and proactive workforce planning.

Michael Watkins
IMD Professor of Leadership and Organizational Change

Managing risk



Identifying pitfalls and
building an ethical and
secure AI



04

Identifying and mitigating potential pitfalls

By Michael Wade and Amit Joshi



Navigating AI risks requires collaboration, transparency, and proactive measures to ensure that the technology works for the benefit of all. By incorporating diverse data, enhancing explainability, developing strong ethical guidelines, investing in human capital, and establishing robust regulatory frameworks, we can steer AI development in a direction that is both innovative and responsible.

Ensuring diverse data and auditing for bias

AI models learn by absorbing information from the data they are trained on. If this data is skewed or incomplete, the model will likely inherit and even amplify these biases, leading to unfair or discriminatory outcomes in areas like hiring, loan applications, and criminal justice. To prevent this, it's crucial to ensure that the data used for training is diverse and representative. This requires diverse data collection and regular bias audits. A good example of an AI auditing tool is Fairlearn, developed by Microsoft, which helps developers build fairer AI systems by offering a dashboard to visualize disparities and algorithms to mitigate them.

Enhancing robustness and security

Embedding security considerations into the AI development lifecycle is crucial for proactively addressing risks and ensuring responsible AI practices. This involves integrating security and ethical checkpoints, impact assessments, and stakeholder consultations throughout the AI development process, from problem formulation to deployment and monitoring.

Adversarial training

Adversarial training involves deliberately exposing an AI model to subtly altered data, called adversarial examples, designed to trick it. These examples are often indistinguishable from normal data to humans, but can cause the AI to make mistakes. By repeatedly exposing the model to these adversarial attacks, it learns to recognize and resist them, improving its accuracy and reliability. This is crucial in any business setting where AI systems are used to make critical decisions, such as loan approvals, risk assessment, or fraud detection.

Testing under real-world conditions

Another essential approach to enhance AI security is to test models under real-world conditions. While controlled lab settings offer a valuable foundation, they often miss the nuances and complexities of real-life scenarios. This can lead to unexpected behaviors and potentially harmful consequences when AI systems are finally deployed. For example, AI systems often interact with complex systems in the real world, like busy traffic networks, fluctuating financial markets, or intricate healthcare systems. These interactions are difficult to replicate accurately in a lab. Real-world testing allows developers to observe how the AI navigates these complex systems and identify any potential issues or unintended consequences.

The human element is another crucial factor. AI systems often interact with people in unpredictable ways. Real-world testing allows developers to observe these interactions and identify any potential usability issues or safety concerns. It's also an opportunity to assess the ethical implications of the AI system in action, identifying any potential biases or discriminatory outcomes.

Keeping humans in the loop

While artificial intelligence holds immense promise, it's crucial to remember that AI systems are ultimately tools shaped by human design and data. This means they can inherit our biases, reflect our flaws, and at times amplify them. To mitigate the risk of AI security issues or ethical lapses, incorporating a human in the loop can be helpful. This means ensuring human oversight and intervention at critical points in the AI system's operation.

Privacy-preserving techniques

As artificial intelligence becomes increasingly intertwined with our lives, the need to protect sensitive information becomes paramount. Privacy-preserving techniques for AI models are emerging as a critical area, addressing the challenge of building AI systems that can learn from data without compromising individual privacy. These techniques, ranging from federated learning to differential privacy and homomorphic encryption, offer innovative solutions for training AI models on sensitive data while keeping personal information secure.

Federated learning

Federated learning is an approach to training AI models that prioritizes privacy. It allows AI systems to learn from data scattered across multiple devices, like smartphones or data centers, without ever needing to directly access or share that raw data. This is a promising approach to privacy preservation, especially in sectors dealing with sensitive information like healthcare or finance.

Differential privacy

Differential privacy is like adding a subtle layer of camouflage to data, making it difficult to identify individuals while still allowing AI models to learn valuable insights. It works by injecting carefully calculated noise into the data or the model's output, ensuring that the presence or absence of any single individual doesn't significantly affect the overall results.

It allows AI models to learn from sensitive data without revealing personal details, opening up new possibilities for AI applications in healthcare, finance, and other privacy-sensitive domains.

Homomorphic encryption

Homomorphic encryption is a new cryptographic technique that allows computations to be performed on encrypted data without ever needing to decrypt it. It's like having a locked box that you can still manipulate and get results from without ever needing to open it. With homomorphic encryption, the manufacturer can encrypt the data before sending it to the AI provider. The AI provider can then perform the analysis on the encrypted data, generate results, and send back the encrypted results. The manufacturer, finally, decrypts the results using their private key, gaining valuable insights about process optimization without ever exposing their confidential manufacturing data.

Managing ethical concerns through AI ethics governance

Managing ethical concerns in AI development and deployment requires robust ethical guidelines. The EU AI Act emphasizes transparency, fairness, and accountability, particularly for high-risk applications like healthcare and law enforcement. Another way to manage ethical concerns is through AI ethics boards.

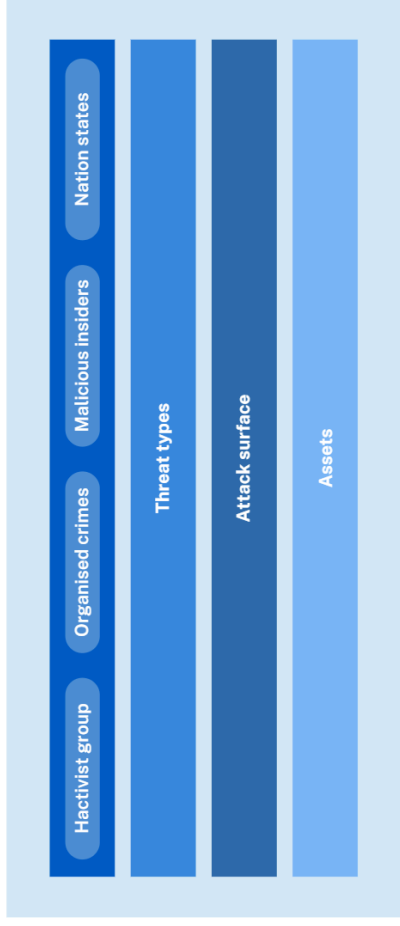


Cybersecurity: How to build a secure-by-design strategy

Scaling AI use broadens the attack surface, making organizations more vulnerable to cyber threats. AI models can be targeted by adversarial attacks, data poisoning, or model inversion techniques that expose sensitive information. The increased connectivity of AI systems creates additional entry points for hackers. Yet, by proactively integrating security measures into AI development and deployment, organizations can minimize risks while maximizing the potential of this tool.

AI introduces unique attack surfaces, requiring proactive and AI-specific security measures. Implementing zero-trust AI, robust data security, and adversarial defenses is critical. It is also crucial to remember that AI security is not just an IT issue; it's a business priority affecting compliance, reputation, and trust.

Threat Actors



Managing risk

Building reliable, accurate, and secure AI requires attention to all four elements of the traditional threat landscape, from actors to types to surface to core assets. AI introduces the prospect of unique attack surfaces within these four threat entry points, requiring proactive and AI-focused security measures.

Secure by design approach

To mitigate risks when building AI processes, follow a secure-by-design approach; you cannot simply fix issues as they come up. A secure-by-design approach integrates security into the design and building of products, systems, and applications, creating AI that is secure by default. A big aspect of this is the AI vendor: if the vendor builds ignoring vulnerabilities, you will bring those into your organization. For example, while there was widespread enthusiasm for the lower cost of DeepSeek, the fast and less expensive Chinese AI software challenger to Silicon Valley stalwarts like OpenAI, Google, or Meta, independent security evaluations of DeepSeek have found a series of weaknesses and exposed susceptibility to cyber threats, including prompt injection attacks, jailbreaking, or data poisoning. Cheaper options may not represent savings in the long term.

AI security strategies

Mitigating cybersecurity risks when using AI requires a multi-layered approach that includes robust security measures, governance frameworks, and continuous monitoring. Here is a checklist for the five key strategies:

- Secure AI model development and deployment
 - Explainability
 - Governance
 - Developer training
 - Secure software development
- Adversarial AI and model manipulation
 - Bias audits
 - Adversarial training
 - Input sanitization
 - Validate training data
- Data privacy and protection
 - Encryption
 - Privacy compliance
 - EU AI Act
 - Access control
- AI supply chain and third-party security
 - Vet external AI tools
 - Security in the cloud
 - Software bill of materials
- AI-specific incident response and monitoring
 - Continuous AI threat monitoring
 - Incident response for AI failures
 - Fair use policies



Building reliable, accurate, and secure AI requires attention to all four elements of the traditional threat landscape, from actors to types to surface to core assets. AI introduces the prospect of unique attack surfaces within these four threat entry points, requiring proactive and AI-focused security measures.

Öykü Işık
IMD Professor of Digital Strategy and Cybersecurity

Takeaways and recommendations

In today's workplace, the ability to interpret and act on data insights directly sets a higher standard for all leaders: management no longer revolves solely around intuition and experience. To extract the most value from data, all leaders should understand the opportunities and risks of AI adoption. As you prepare to set your organization's AI strategy, remember that this is not a straightforward systems upgrade; it's a transformation of how you do business, and it is an ongoing process. At the heart of the transformation is the cultural shift, ensuring that everyone in your organization, from top leadership and across the entire workforce, understands the value-data-people framework. Keep the competitive advantage this brings to your organization – and to your people – top of mind.

Measuring the impact of your transformation is essential. Successful organizations systematically track outcomes. Any major AI investment should include a clear plan and budget for ongoing evaluation. With the right strategy and a commitment to measuring impact, AI can become a powerful driver of productivity, innovation, and long-term growth. As you prepare to lead your organization's transformation, keep in mind this north star: we must leverage the strengths of both AI and human intelligence. AI can process vast amounts of data and identify patterns that humans might miss, while humans bring critical thinking, empathy, and contextual understanding to the project. By combining these strengths and ensuring meaningful human oversight, we can harness the power of AI while mitigating its risks, paving the way for a more ethical and secure future in which AI empowers individuals, organizations, and society as a whole, while safeguarding against unintended consequences.

Four questions every leader should ask

1. Have we built our strategy around the value-data-people framework?
2. Are we leveraging the best of artificial and human intelligence, together?
3. Can we strengthen our understanding of how GenAI is transforming our industry and the different functions across our organization?
4. Is our AI strategy secure by design?








Learn more about the research at the TONOMUS Global Center for Digital and AI Transformation.

With contributions from:

- Julia Binder**
Professor of Sustainable Innovation and Business Transformation
- Didier Bonnet**
Professor of Strategy and Digital Transformation
- Mark Greeven**
Professor of Management Innovation and Dean of Asia
- Öykü Işık**
Professor of Digital Strategy and Cybersecurity
- Amit Joshi**
Professor of AI, Analytics, and Marketing Strategy
- José Parra Moyano**
Professor of Digital Strategy
- Misiek Piskorski**
Professor of Digital Strategy, Analytics and Innovation and Dean of Executive Education
- Michael Wade**
TONOMUS Professor of Strategy and Digital, and Director of the TONOMUS Global Center for Digital and AI Transformation
- Michael D Watkins**
Professor of Leadership and Organizational Change
- Howard Yu**
LEGO® Professor of Management and Innovation

Build the skills to lead in a digital world
Learn with IMD: explore our Digital transformation and AI programs.

Who for	Outcome	Learning objectives	Delivery
Executives and leaders leading digital transformation Professionals looking to embed digital and AI in their business	Mastery of AI-driven transformation 	Master the core skills, get the right tools, frameworks, language, and capabilities to drive change	5 days On campus Lausanne and Riyadh
Business leaders and managers with a foundational understanding of AI and analytics, looking to dive deeper into the topic and set strategy	Proficiency in AI strategy and execution 	Learn to design and implement practical AI strategies tailored to your business	5 days On campus Lausanne
Leaders and entrepreneurs in all business areas looking to leverage new technologies to drive fast innovation and growth Ideal for individuals and teams	Hands-on innovation skills 	Explore how to apply AI and new technologies for innovation and business growth	5 days On campus Lausanne
Leaders driving digital and AI transformation in a rapidly evolving business environment	Immersive, practical learning to boost AI expertise 	Elevate your business impact, bridging gaps in digital and AI knowledge with actionable insights	5 days On campus Singapore
Professionals looking to upskill and advance their career with the recent advancement in applications of AI for business	Rapid skills development 	Rapidly learn how to apply generative AI to your business or personal life	1 week Online

