Making AI Work for Business

Introduction

Artificial intelligence is creating significant opportunities for companies to improve efficiency, enhance decision-making, develop new products and services, and fundamentally change how work is performed.

Yet there is an important distinction between experimenting with AI and establishing the capabilities needed to use it effectively across a company.

Many businesses have started by introducing AI assistants, testing individual applications or automating specific tasks. These initiatives can deliver useful results, but deploying more AI tools does not necessarily translate into greater business value.

The challenge is to move beyond isolated applications and develop the information, technology, processes and controls that allow AI to operate effectively within the business.

The objective is not AI adoption itself, but measurable improvements in business performance, whether through increased productivity, lower operating costs, improved customer service, better decision-making or new opportunities for growth.

One way of understanding this challenge is through a concept inspired by psychologist Abraham Maslow's Hierarchy of Needs.

Maslow proposed that human needs could be understood as a hierarchy, ranging from fundamental requirements to higher aspirations. A similar principle can be applied to AI: more advanced applications depend on a combination of underlying capabilities that enable them to operate reliably, securely and at scale.

This provides the basis for an **AI Hierarchy of Needs**, comprising five interconnected levels: Foundations, Control, Integration, Scale and Transformation.

1. Foundations: Giving AI the Information and Capabilities It Needs

The first level concerns the information, knowledge and technology that AI requires to perform useful work.

Every company has accumulated considerable business knowledge and intellectual property. This includes documents, datasets, methodologies, business rules, operational procedures, processes and information held within existing systems.

Much of this information is distributed across different applications, business functions and data repositories. Some is formally documented, while other knowledge is embedded in business processes and the experience of employees.

For AI to perform meaningful business activities, it needs access to the relevant information and an understanding of how that information should be used.

For example, an AI application supporting pension administration may need access to scheme documentation, member records, calculation methodologies, regulatory requirements and established administrative processes and procedures.

Access to information alone is not sufficient. The application also needs to understand which information is relevant, how different sources relate to one another and which rules govern its use.

This does not mean that companies must consolidate all their information into a single system before adopting AI. Existing information and systems can remain in place, provided the necessary knowledge, data and capabilities can be made accessible in a structured and controlled way.

The objective is to establish a foundation that allows AI to operate within the context of the business rather than relying solely on the large language model.

2. Control: Establishing Trust, Governance and Accountability

The second level concerns the controls needed to deploy AI responsibly and reliably.

AI applications operate within existing business environments, where activities are subject to rules, permissions, regulatory obligations and established responsibilities.

Introducing AI does not remove these requirements.

A company must be able to determine what information an AI application can access, what actions it is permitted to perform, which decisions require validation and when human approval is necessary.

Some activities can be performed using AI models, while others require deterministic rules or calculations that produce consistent, verifiable results.

For example, AI might interpret incoming correspondence and identify the action required, but a financial calculation may need to follow prescribed business rules rather than relying on a probabilistic model.

Effective governance therefore involves more than monitoring model outputs. It extends to the management of agents, business tools, workflows, data access, approvals, audit trails and operational performance.

Companies also need visibility over AI usage and associated costs, particularly where applications involve multiple models, external services or significant transaction volumes.

These controls should form part of the design and implementation of AI applications from the outset.

The objective is not to restrict innovation, but to create the framework and conditions in which AI can be trusted to undertake increasingly important business activities.

3. Integration: Moving Beyond Individual AI Applications

The third level is where AI begins to move from supporting individual tasks to operating within complete business processes.

Many early AI applications focus on relatively self-contained activities, such as summarising documents, drafting correspondence, taking minutes, extracting information or answering questions.

These capabilities are valuable, but business processes rarely consist of a single activity.

A customer onboarding process, for example, may involve collecting information, validating documents, checking eligibility, applying business rules, updating systems, generating correspondence and obtaining approvals.

Delivering this process requires more than an individual AI agent.

It involves coordinating AI agents with existing applications, business tools, deterministic rules, data sources and human activities through an integrated workflow.

Different activities may require different technologies. Large language models may be appropriate for interpreting unstructured information, while conventional software, rules engines or existing systems may be better suited to calculations, validations and transactional processing.

The important consideration is how these capabilities work together.

This is also why AI should not be viewed as a replacement for every existing technology. In many cases, the greatest value comes from combining AI with the systems and capabilities a company already has.

Integration enables AI to contribute to end-to-end business processes, rather than remaining confined to isolated tasks.

4. Scale: Building Reusable Capabilities Across the Company

The fourth level concerns the ability to extend AI adoption without repeatedly developing the same capabilities.

A company may initially develop an AI application to address a particular requirement. In doing so, it may create agents, business tools, integrations, workflows, rules and controls that could also be used for other processes.

If these components are designed for reuse, subsequent applications can be developed more efficiently.

For example, a tool created to retrieve customer information may be useful in onboarding, customer servicing, complaints handling and other processes. Similarly, document interpretation, validation, approval and reporting capabilities may support multiple applications.

Over time, these components can form a growing collection of reusable business capabilities.

This changes the economics of AI adoption.

Rather than treating each new application as a separate development exercise, companies can build upon previous investments, reducing duplication and accelerating implementation.

It also supports consistency in governance, security and operational management.

However, scaling AI involves more than reusing technology. Companies must also consider how applications are deployed, monitored, maintained and managed across different business functions and activities.

The objective is to establish a company-wide capability that supports the development and operation of multiple AI-enabled solutions while retaining the flexibility to address specific requirements.

5. Transformation: Rethinking How Work Is Performed

The highest level of the hierarchy concerns the opportunities that arise when AI becomes an established part of how a company operates.

Much of the initial focus on AI has been on improving existing activities: completing tasks faster, reducing manual effort, improving accuracy and lowering costs.

These benefits can be substantial.

However, the longer-term opportunity extends beyond making existing processes more efficient.

Rather than automating individual stages of a business process, a company could redesign the process around AI agents, intelligent workflows and existing systems, eliminating unnecessary steps and handovers, reducing manual intervention and enabling employees to focus on activities requiring expertise and judgement.

The same principles can support the development of new products, services and business models.

Transformation therefore involves more than adopting new technology. It requires companies to examine their existing operating models and identify opportunities to perform work in fundamentally different ways.

Importantly, transformation does not mean removing human involvement. It means determining where human expertise, judgement and accountability add the greatest value and designing processes accordingly.

6. The Hierarchy Is Not a Rigid Sequence

Although the five levels form a hierarchy, they should not be interpreted as a series of stages that must be completed before moving to the next.

Companies do not need to establish an extensive AI infrastructure or reorganise all their business information before developing their first useful application.

Nor should governance, integration and scalability be treated as activities that begin only after an initial implementation has been completed.

A more practical approach is to start with a specific business requirement and develop the capabilities needed to address it.

This might involve identifying relevant information, connecting existing systems, implementing business rules, developing agents and workflows, and establishing the necessary controls.

Once the application has been implemented and its value demonstrated, the components developed can be retained and reused.

Further applications can then build upon these capabilities.

In this way, the hierarchy describes **interdependent requirements rather than a prescribed implementation sequence**.

Each application can deliver immediate business value while contributing to a growing foundation for wider AI adoption.

This approach also reduces the risk of investing heavily in technology and infrastructure before establishing where AI can deliver measurable benefits.

7. Assessing AI Maturity

The hierarchy also provides a useful framework for assessing how effectively a company is adopting AI.

Rather than measuring progress by the number of AI applications or agents deployed, companies can examine the capabilities they have established across the five levels.

Five questions provide a starting point:

Foundations: Can AI applications access and use the relevant business information, knowledge, data and systems?

Control: Are AI activities governed by appropriate security, business rules, validation, monitoring and accountability?

Integration: Can AI agents and tools operate within existing systems and end-to-end business processes?

Scale: Can the components developed for individual applications be reused, managed and deployed across the company?

Transformation: Is AI enabling the company to rethink how work is performed and develop new capabilities, products or services?

Progress should therefore be assessed across all five levels, recognising that weaknesses in foundational capabilities, governance or integration can limit the benefits of more advanced AI applications.

A company may be advanced in some areas and less developed in others. The purpose is to identify strengths, gaps and priorities for further investment.

A relatively small number of well-integrated, governed applications may represent greater progress than a large collection of disconnected AI tools.

Ultimately, AI maturity should be assessed by the business value being delivered and the company's ability to extend that value sustainably.

8. Building AI Capability One Application at a Time

The AI Hierarchy of Needs highlights an important distinction between deploying individual AI applications and developing the capabilities needed to use AI effectively across a company.

Individual applications can deliver immediate benefits, but their longer-term value increases when they contribute to a growing collection of reusable business knowledge, agents, tools, workflows, integrations and controls.

This requires an approach that combines practical implementation with a longer-term view of how AI capabilities can develop.

Companies should begin with clearly defined business requirements, establish the information and controls needed to address them, and develop solutions that can be extended and reused.

As these capabilities accumulate, the cost and complexity of subsequent implementations can be reduced, while the opportunities for integrating AI into wider business processes increase.

The objective is to establish the capabilities needed to use AI effectively, securely and consistently across the company, delivering measurable improvements in business performance while creating opportunities for growth and transformation.

About Kognoesis

Kognoesis is an Agentic AI platform for building, deploying and governing AI agents and intelligent workflows. It provides the development, deployment and management capabilities to create AI-enabled solutions and applications specific to a company's requirements, while supporting their integration, governance and reuse across business functions.

The platform addresses the five levels of the AI Hierarchy of Needs:

- **Foundations:** The Kognoesis Business Repository provides structured access to business information and knowledge, including documents, datasets, methodologies and business rules. Integration capabilities enable AI applications to access and work with information held in existing systems, providing the business context needed to perform meaningful activities.
- **Control:** Kognoesis incorporates governance and management controls covering AI agents, tools and workflows, including access permissions, deterministic business rules, validation, human approvals, monitoring and token usage and cost tracking. Its LLM-agnostic architecture, including support for local models, provides flexibility in addressing data security, privacy and deployment requirements.
- **Integration:** Kognoesis provides the development capabilities to build, configure and deploy AI agents, reusable business tools, deterministic rules, integrations and intelligent workflows. These capabilities can be combined with existing business systems and processes to develop integrated, end-to-end AI-enabled solutions rather than relying on isolated agents or applications.
- **Scale:** Agents, business tools, integrations, workflows and controls developed using Kognoesis can be reused and extended across subsequent implementations. This supports faster development, reduces duplication and enables companies to progressively extend AI adoption across business functions and processes.
- **Transformation:** By combining these capabilities, Kognoesis enables companies to move beyond automating individual tasks towards redesigning business processes and developing new AI-enabled products, services and ways of operating.

Kognoesis does not require companies to complete each level of the hierarchy before progressing to the next.

Companies can begin with a specific business requirement, develop an application that delivers measurable value and establish the necessary information, integration and governance capabilities as part of that implementation.

These capabilities can then be extended and reused in subsequent applications, creating a growing foundation for company-wide AI adoption.

The Kognoesis platform enables companies to develop practical AI solutions around specific business requirements, while progressively building the reusable capabilities needed to develop, deploy, integrate, manage and govern AI at scale.

Kognoesis