

Knowing the Moves Isn't the Same as Knowing How to Play

What chess can teach us about implementing AI across an organisation

Learning chess starts with learning the rules. You learn how each piece moves, how pieces are captured, how to castle and what constitutes checkmate.

But knowing how the pieces move does not mean you necessarily know how to play chess.

Effective play requires an understanding of the position, the objective, the available options, the threats and opportunities and the likely consequences of each move. You need tactics, but you also need strategy.

There is an interesting parallel with the way organisations are beginning to implement AI.

AI models can already perform a range of individual tasks. They can interpret documents, analyse information, generate content, interrogate data, make recommendations and interact with other systems. Agentic AI extends this further by enabling AI to use tools and take actions.

But being able to perform individual tasks is not the same as delivering tangible business benefits such as improved productivity, better customer outcomes, reduced operational costs, improved service and increased business capacity.

From moves to strategy

A useful discipline when learning chess is to ask three questions before making a move:

What is my opponent threatening?

What checks, captures or threats are available to me?

If I make my intended move, what is their strongest reply?

A move should not be considered in isolation. It needs to be understood and made in the context of the position and what might happen next.

The same applies to developing, implementing and operating AI-enabled solutions and agents within a business. They need to be considered in the context of the organisation, how it operates and what it is seeking to achieve.

It requires context.

The organisation provides the context

That context already exists within organisations.

It is contained in their intellectual property, knowledge, methodologies, business rules, data, processes and technology. Some of it is formally documented. Some is embedded in systems, procedures and controls. And some resides in the experience and judgement of the managers and teams who perform the tasks.

Together, these are fundamental to what makes an organisation distinctive and enables it to compete effectively in its market. They represent knowledge, expertise and capabilities that have often been developed and refined over many years and which competitors cannot necessarily replicate.

This is where an important distinction emerges. The underlying AI models may increasingly be available to everyone, but the knowledge, intellectual property, methodologies, business rules, data, processes and technology with which they are combined are specific to each organisation.

Organisations can therefore use AI to build on these existing strengths and develop solutions specific to their business.

As access to increasingly capable AI models becomes widespread, the opportunity lies not in access to the models themselves, but in how effectively organisations combine and implement them with the proprietary knowledge, capabilities and intellectual property that already underpin their competitive advantage.

Tactics are not strategy

Chess makes an important distinction between tactics and strategy.

A tactic addresses an immediate opportunity: winning a piece, creating a threat or delivering checkmate.

Strategy considers the wider position and asks what needs to be achieved over a longer sequence of moves to achieve a victory.

Much of the first wave of AI implementation has understandably been tactical.

Can AI extract information from this document?

Can it summarise this case?

Can it draft this communication?

Can it analyse this dataset?

Can an agent automate this particular task?

Each can produce worthwhile business benefits.

But as organisations implement AI, a broader question emerges:

How do these individual capabilities fit into the wider processes through which the organisation operates and delivers its services, and how can they be combined into solutions that deliver meaningful business outcomes?

A customer onboarding process, for example, may involve receiving and interpreting information, validating it against business rules, identifying missing information, communicating with the customer, accessing external data, updating multiple systems, obtaining approvals and managing exceptions.

AI may assist with many of these activities individually. The greater opportunity comes from bringing them together within an AI-enabled solution that can support and automate more of the end-to-end process.

That requires orchestration across AI models and agents, deterministic business rules, existing systems, data, people and work management.

Operating within defined boundaries

The rules of chess define what a player can and cannot do. Within those rules, there are an enormous number of possible moves and positions, and deciding what to do requires strategy and judgement.

There is a useful parallel when developing AI-enabled solutions. AI models and agents can reason and determine what action to take, but they similarly need to operate within defined boundaries. These may be determined by business rules, data, systems, processes, controls and where human judgement or approval is required.

AI reasoning therefore needs to be combined with deterministic business rules, permissions, controls and guardrails appropriate to the process and the organisation. This enables AI models and agents to operate within clearly defined parameters, while retaining human judgement and approval where required.

The technology used also needs to reflect the nature of the information being processed.

Where processes involve personal, confidential or commercially sensitive information, organisations may require that information to remain within their own controlled environment. The ability to use locally deployed AI models alongside, or instead of, external frontier models provides an important option for these use cases.

The question is therefore not simply:

What can the AI do?

It is:

How should AI be used, what should it be allowed to do, using what information, under what circumstances and subject to what controls?

Strong players recognise patterns

Strong chess players do not calculate every position entirely from first principles. Experience allows them to recognise patterns.

They have encountered similar structures, attacks and defensive positions before. That accumulated knowledge helps them identify promising moves and concentrate their analysis on what matters.

Organisations implementing AI should be able to build capability in much the same way.

Knowledge, documents and data used for one solution may also be relevant to another. An integration built to access a core system can be reused. A business tool created for one process can be available to others. Business rules, knowledge sources, data capabilities, AI agents and other components can similarly become reusable building blocks.

The first AI-enabled solution therefore does more than solve one problem.

Done correctly, it begins to create an organisational capability.

The second solution can reuse elements of the first. The third can reuse elements of both. Over time, the organisation progressively develops a library of knowledge, agents, business tools, integrations, rules and data capabilities that can be combined in different ways to address different requirements.

The value of an AI platform should compound as it is used.

Building capability one move at a time

This does not mean organisations need to design an enterprise-wide AI architecture before implementing anything.

Chess games are still played one move at a time.

Similarly, the practical approach to implementing AI is often to start with individual use cases where there is a clear business problem and demonstrable value.

The important difference is to implement those use cases with the broader objective in mind.

Rather than creating isolated AI applications, organisations can develop AI-enabled solutions around their specific requirements, while progressively developing and reusing the

knowledge, agents, business tools, integrations, rules and controls created for individual use cases.

Each implementation can therefore contribute to a growing set of capabilities that can be applied to subsequent use cases and extended across different processes.

This approach also provides consistency in how solutions are developed, implemented, operated, governed and controlled as the use of AI expands across the organisation.

Over time, individual implementations begin to connect. AI moves from performing isolated tasks to supporting and automating end-to-end processes and, progressively, working across processes and systems throughout the organisation.

From individual solutions to organisational capability

This is the thinking behind Kognoesis.

Kognoesis is an enterprise Agentic AI platform through which organisations can develop, implement and govern AI-enabled solutions and applications specific to their requirements.

The platform is designed to adapt to the organisation, rather than requiring the organisation to adapt the way it works to fit the platform. This includes the flexibility to work collaboratively with clients to enhance and extend the platform where required to meet their specific requirements.

Solutions and applications developed within Kognoesis can draw on the organisation's existing intellectual property, knowledge, methodologies, business rules, data, processes and technology, with the platform's Business Repository providing structured access to the business information, knowledge and data required by AI models and agents. Kognoesis combines AI models and agents with the organisation's knowledge and data, workflow, deterministic business rules, reusable business tools, integrations, existing systems and human judgement, enabling organisations to develop AI-enabled solutions and applications that support and automate end-to-end business processes.

The business tools and other capabilities developed within these workflows can be reused across other workflows and use cases. Knowledge, agents, integrations, rules and data capabilities can similarly be reused and progressively extended, enabling each new solution to build on capabilities already developed within the platform.

Kognoesis is LLM-agnostic, allowing different models to be selected according to the requirements of each solution, including locally deployed models where personal, confidential or commercially sensitive information needs to remain within the organisation's environment.

In this way, organisations can start with individual solutions that deliver tangible business benefits, while progressively developing a broader organisational capability to use AI effectively and extend it across the business.

Learning to play the game

When you first learn chess, you learn the rules and how the individual pieces move.

But learning to play the game effectively requires much more. It requires understanding the position, recognising threats and opportunities, developing a strategy and considering the consequences of each move.

There is a similar transition in the way organisations are approaching AI.

The question is moving from:

“What can AI do?”

to:

“How can we use AI to build and improve on what makes our organisation distinctive and the way our business actually works?”

That means effectively combining AI with the organisation's own knowledge, intellectual property, business rules, data, processes and technology, within the appropriate controls.

The objective is not simply to use more AI.

It is to develop the organisational capability to use AI effectively.