

AI Approaches in Business: A CIO's Decision Making Guide

*Understanding Rule-Based, GenAI, and Agentic AI
Through Real-World Outcomes*

85%

of AI pilots fail to reach production.
Here's what actually works.

The AI hype cycle has been fascinating to watch. After a year of breathless excitement about agentic AI and autonomous systems, McKinsey's recent research confirms what we've been telling clients all along: **there's no one-size-fits-all AI solution**. Success comes from understanding when to use which approach and, critically, when to use none at all.

At SynaptyX, we've implemented everything from basic rules engines processing millions of transactions to sophisticated multi-agent systems coordinating across enterprise workflows. This guide shares what we've learned, backed by McKinsey's latest findings, about choosing the right AI approach for your actual business needs.



Understanding the Three Approaches

Let's cut through the marketing speak. Here's what each AI approach actually means in practice:

What They Actually Do

RULE-BASED AI/ML

- ▶ Executes predefined logic and learned patterns
- ▶ Processes structured data with consistency
- ▶ Makes decisions based on explicit rules
- ▶ Zero ambiguity, maximum predictability

GENAI (LLMS)

- ▶ Understands and generates human-like text
- ▶ Handles unstructured, ambiguous inputs
- ▶ Provides reasoning and explanations
- ▶ Context-dependent, probabilistic outputs

AGENTIC AI

- ▶ Coordinates multiple AI systems and tools
- ▶ Plans and executes multi-step workflows
- ▶ Adapts approach based on outcomes
- ▶ Maintains state across complex processes

Best Used For

RULE-BASED AI/ML

- ▶ High-volume transactional processing
- ▶ Regulatory compliance and audit trails
- ▶ Fraud detection and anomaly identification
- ▶ Credit scoring and risk assessment
- ▶ Quality control and validation

GENAI (LLMS)

- ▶ Document analysis and summarisation
- ▶ Customer communication and support
- ▶ Content creation and translation
- ▶ Code generation and review
- ▶ Research and insight extraction

AGENTIC AI

- ▶ End-to-end process automation
- ▶ Cross-system workflow orchestration
- ▶ Complex decision-making chains
- ▶ Investigative and analytical tasks
- ▶ Adaptive customer journeys

Key Characteristics

RULE-BASED AI/ML

- ▶ Deterministic and explainable
- ▶ Fast and cost-efficient at scale
- ▶ Requires explicit programming
- ▶ Struggles with novel scenarios

GENAI (LLMS)

- ▶ Handles nuance and ambiguity
- ▶ Requires careful prompt engineering
- ▶ Potential for hallucinations
- ▶ Context window limitations

AGENTIC AI

- ▶ Breaks goals into executable steps
- ▶ Maintains state across processes
- ▶ Requires robust guardrails
- ▶ Exponential "Context Tax"

Commercial Reality

RULE-BASED AI/ML

Often the most cost-effective solution. A well-designed rules engine at £10K can outperform a £100K LLM implementation for structured decisions. Processing millions of transactions? This is your workhorse.

GENAI (LLMS)

At £0.03 per 1K tokens, a single complex conversation can cost pounds. Brilliant for high-value, low-volume tasks (contract analysis, executive summaries). Expensive disaster for high-volume, low-value ones.

AGENTIC AI

McKinsey found that "workflow centricity" determines success. The fancy agent means nothing if the underlying workflow is broken. Most "agentic" use cases perform better as well-orchestrated traditional automation.

The Context Tax: Hidden Costs

Every AI approach carries hidden costs we call the "Context Tax":

- ▶ **Rule-based AI:** Zero context tax but potentially infinite edge cases to handle
- ▶ **GenAI:** Linear cost growth, maintaining conversation state across thousands of interactions becomes expensive quickly
- ▶ **Agentic AI:** Exponential cost growth, each agent maintaining context while coordinating with others

Example: Processing 10,000 customer enquiries monthly

- ▶ Rules engine: ~£500/month infrastructure
- ▶ GenAI with full context: ~£15,000/month in tokens alone
- ▶ Agentic coordination: ~£45,000/month plus orchestration overhead

When to Use Which Approach

Use Rule-Based AI/ML when:

- ✓ You can define clear if-then logic
- ✓ Consistency and auditability are paramount
- ✓ Processing volume > 10,000 transactions/month
- ✓ Accuracy requirements > 99%
- ✓ Regulatory compliance is critical

Use GenAI/LLMs when:

- ✓ Dealing with unstructured text or natural language
- ✓ Need to explain complex information simply
- ✓ Creating content or summaries for human review
- ✓ Handling nuanced interpretation or sentiment
- ✓ Value of insight > cost per interaction

Use Agentic AI when:

- ✓ Coordinating across 3+ systems or data sources
- ✓ Workflow has multiple decision branches
- ✓ Value of automation > £100K annually
- ✓ Human oversight can handle edge cases
- ✓ Process improvements compound over time

Combine approaches when:

Most real-world solutions are hybrid. As McKinsey notes, successful implementations use "analytical AI components for quantitative analysis alongside GenAI for natural-language tasks."

Uncomfortable Truths About AI

LLMs are poor calculators: For mathematical operations, a simple spreadsheet outperforms the most sophisticated language model. Don't use GenAI for what Excel does perfectly.

Context windows are a constraint: Every piece of context costs money and processing time. Most context is noise, focus on what's essential.

Hallucination isn't always bad: Creative interpretation can be valuable (marketing copy) or catastrophic (compliance documents). Know where you need creativity versus accuracy.

Agents aren't magic: McKinsey found the successful agents were those solving specific, bounded problems, not general-purpose "AI employees."

Real Stories from the Field



Digital Marketing: Manual campaign analysis across platforms was causing delays. Solution: secure data pipelines with custom metrics achieved 1-3% CAC reduction and 10% efficiency gains. Lesson: Sometimes better analytics beats "intelligent" automation.

Financial Digitisation: Processing 4,500 cities' financial statements from scanned PDFs. Solution: OCR for extraction + LLM for interpretation resulted in 85% effort reduction. Lesson: Combining approaches outperforms single solutions.

Build vs Buy vs Ignore Decision

Ignore AI if:

- ✓ Process variance < 20%
- ✓ Volume < 1K/month
- ✓ Human cost < £50K/year
- ✓ No competitive disadvantage
- ✓ ROI timeline > 18 months

Buy AI if:

- ✓ Solved problem (OCR, sentiment)
- ✓ Vendor solutions mature
- ✓ Integration < 3 months
- ✓ Total cost < building

Build AI if:

- ✓ Core competitive differentiation
- ✓ Proprietary data advantage
- ✓ 10x improvement potential
- ✓ Impact > £1M annually

The SynaptyX Approach

Our philosophy, informed by both experience and McKinsey's findings:

1. **Start with the work to be done**, not the technology available
2. **Map the complete workflow**, understanding human and system touchpoints
3. **Identify automation opportunities**, but keep humans for judgment calls
4. **Choose the right tool for each step**, rules, GenAI, or agents as appropriate
5. **Measure actual business outcomes**, not technical metrics

As McKinsey observed: "Organisations that focus too much on the agent itself often end up with great- looking agents that don't actually improve the overall workflow."

Three Principles for Success

1

Respect the Economics

Rules engines for volume. LLMs for value. Agents for complexity. Choose based on economics, not excitement.

2

Perfect the Workflow First

"Brilliant AI deployed badly is just badly deployed AI." Fix your processes before adding intelligence.

3

Keep Humans in the Loop

People oversee accuracy, ensure compliance, use judgment, and handle edge cases. Design for collaboration, not replacement.

Practical Implementation

In typical end-to-end processes (onboarding customers, processing claims, investigating cases), you don't pick one AI approach. You design the workflow, then assign the right tool to each step:

- ▶ **Rule-Based AI/ML** for clear decisions, calculations, and compliance checks
- ▶ **GenAI** for reading, explaining, and drafting
- ▶ **Agentic AI** for orchestrating steps and managing the overall workflow

The best implementations, as McKinsey found, achieve "superior integration and efficiency" by customising combinations rather than forcing off-the-shelf solutions.

The Bottom Line

After a year of agentic AI reality checks, the industry is learning what we've believed from the start: success isn't about having the most sophisticated AI. It's about:

- ▶ Understanding your actual workflows
- ▶ Choosing the right tool for each task
- ▶ Measuring real business outcomes
- ▶ Building with your people, not despite them

**The future isn't agents everywhere.
It's the right tool, in the right place, working with the right people.**

At SynptyX, we don't start with AI. We start with your business and your people. Our Lattice accelerators aren't finished products you're meant to squeeze into, they're starting points you can shape to fit your reality. Build fast. but build right.

Contact Us



corp@synptyx.ai



www.synptyx.ai