ARTIFICIAL CONTROL

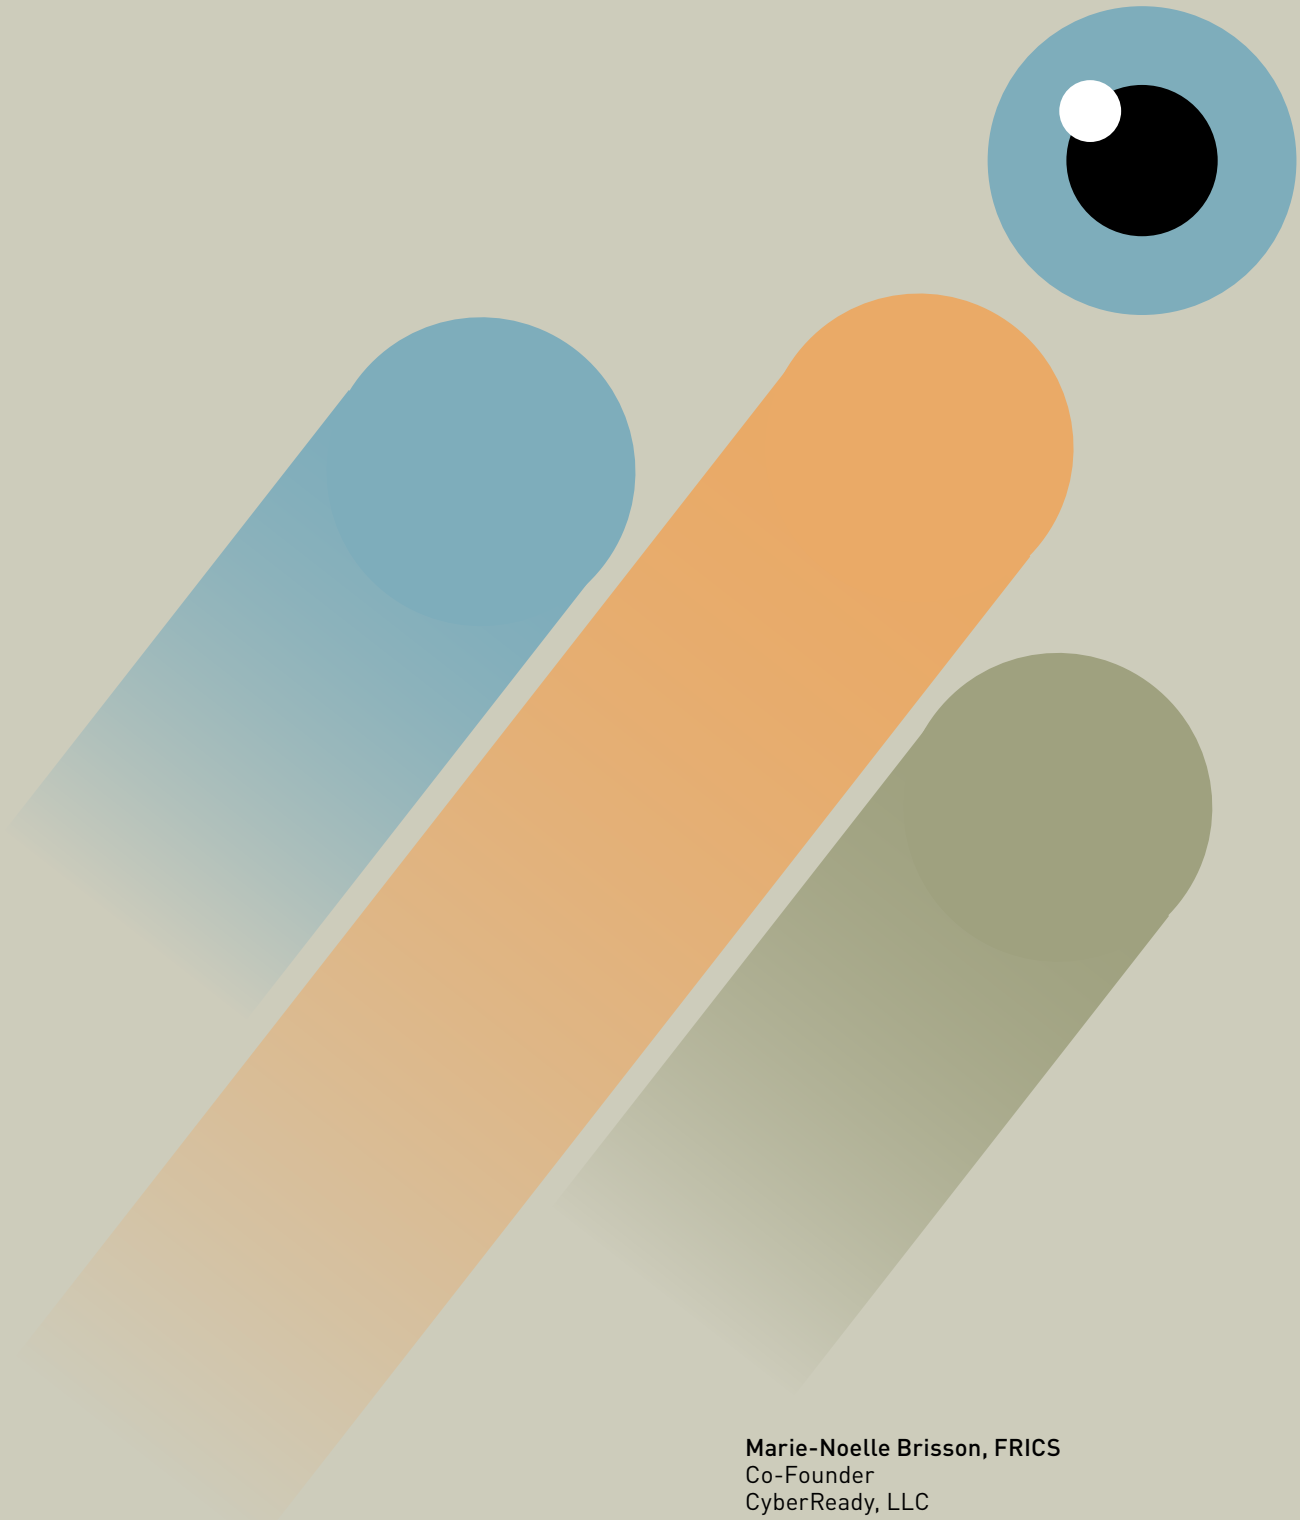

Marie-Noelle Brisson, FRICS
Co-Founder
CyberReady, LLC

Michael Savoie, PhD
Co-Founder
CyberReady, LLC

For CRE professionals, AI might seem like another tool in the smart building arsenal. But the possibilities of AI go far beyond utility and will require smart management and oversight to maximize its potential (and avoid its pitfalls).

These days, artificial intelligence (AI) as a term represents the emergent age of intelligent machinery—far beyond robotics and science fiction. However, the concern for use, management, and oversight of this new technology remains the same as it always has for revolutionary technologies: how do we harness the positive capabilities of the technology without incurring a loss of control?

For CRE professionals, AI might seem like another tool in the smart building arsenal. But the possibilities of AI go far beyond utility and will require smart human management and oversight to maximize its potential (and avoid its pitfalls).

DEFINING AI

AI broadly refers to computer systems that can perform complex tasks normally done by human-reasoning, decision making, creating, and so forth. Because AI is still being realized and developed, it does not yet have a single, simple definition, especially because its tools are capable of a wide range of tasks and outputs.[1]

For the purposes of this paper, we are focusing on AI that excels at specialized tasks, such as:

- Image recognition: Identifying objects or features within images.
- Speech recognition: Converting spoken language into text.
- Natural language processing: Understanding and responding to human language.
- Recommendation systems: Suggesting content based on user preferences.
- Virtual assistants: Providing information and performing tasks based on user commands (e.g., Siri, Alexa, etc.).

Key Characteristics of this type of AI include specialized design – usually for a single purpose, reactive response to specific inputs, and lack of self-awareness. The two most common AI applications today are Generative AI (GenAI) and Agentic AI.[2]

GenAI is a type of artificial intelligence that focuses on creating new and original content, such as text, images, music, videos, and code. GenAI models learn from vast amounts of data (i.e., large language models [LLMs]) and then use that knowledge to generate novel outputs that resemble the characteristics of the training data.

Agentic AI, or AI agents, refers to AI systems capable of independent decision-making and autonomous behavior. These systems can reason, plan, and perform actions, adapting in real time to achieve specific goals. Examples include AI-powered robots, autonomous vehicles, and sophisticated customer service agents.

Limitations of these types of AI include lack of general intelligence; they cannot transfer knowledge or skills from one task to another; limited adaptability, and potential for bias. This last limitation is, and should be, of key concern to organizations intending to utilize AI tools.

AI tools "learn" from the large language models (LLMs) provided to them. If the LLM is created from a specific perspective, then the AI tool will have that perspective. Think about CRE. If the LLM is built around the premise that single family home sales are proportional to interest rates, then the AI tool will base it decision on rising interest rates result in greater sales. The AI tool only knows what's in the LLM. The data must be objective and as unbiased as possible if the AI is to provide accurate analysis. A great example of this was the flawed initial roll out of Google Gemini.[3]

EXAMPLES AND OPPORTUNITIES FOR REAL ESTATE PROFESSIONALS

As is well-established, AI is already helping CRE professionals process large volumes of market research data, automate repetitive tasks, and generate insights that inform strategic decisions in many trades.

AI systems can continuously monitor data handling practices to ensure personal information is processed according to privacy law requirements. This includes automated data classification, consent management, and breach detection - critical capabilities given that real estate transactions involve extensive personal and financial data over rather long holding periods compared to consumer dealings.

Beyond those common uses, AI can produce unintentional benefits such as enhancing customer or tenant experience, better understand climate risks or breaking silos within an organization to name a few.

AI capabilities can be integrated into existing BIMs: for instance, by combining BIM floor plans with occupancy sensors and access card data, AI can provide detailed insights into how tenants actually use space and make facility managers more responsive. Beyond this obvious operational benefit, upstream, AI also eases the relationship to asset managers: it can track tenant satisfaction metrics, predict renewal likelihood, and identify opportunities to improve tenant retention. Today, there are many machine-learning software programs offering plug-and-play solutions (i.e., Elise[AI], Proda[AI], Prophia, LeaseLens, Pipe.CRE, etc.).

Machine learning algorithms automate many underwriting tasks help detect fraud and facilitate compliance reviews; but can also enhance the understanding of climate risks at the property level and across portfolios, thus helping with site selection, building design or rehabilitation, insurance pricing, and capital expenditures to name a few.

Another positive collateral of AI is that its pervasive use encourages data and information collection and conversations across business units. As illustrated in *Exhibit 1,* digital risks are often handled in a fragmented way, where each department only focuses on its specific threats without understanding the broader risk landscape.[4]

Aligning teams and removing silos is essential in establishing and maintaining strong communication and collaboration to increase digital risk management effectiveness. It is a necessary first step in designing a holistic risk-handling culture, but requires commitment to change management. AI could be an agent of change, for example encouraging conversations between the CISO, IT manager, facility manager and compliance officer in designing use cases for predictive maintenance.

EXHIBIT 1: SILOED RISK: WHO IS RESPONSIBLE FOR MANAGING DIGITAL RISK IN YOUR ORGANIZATION?

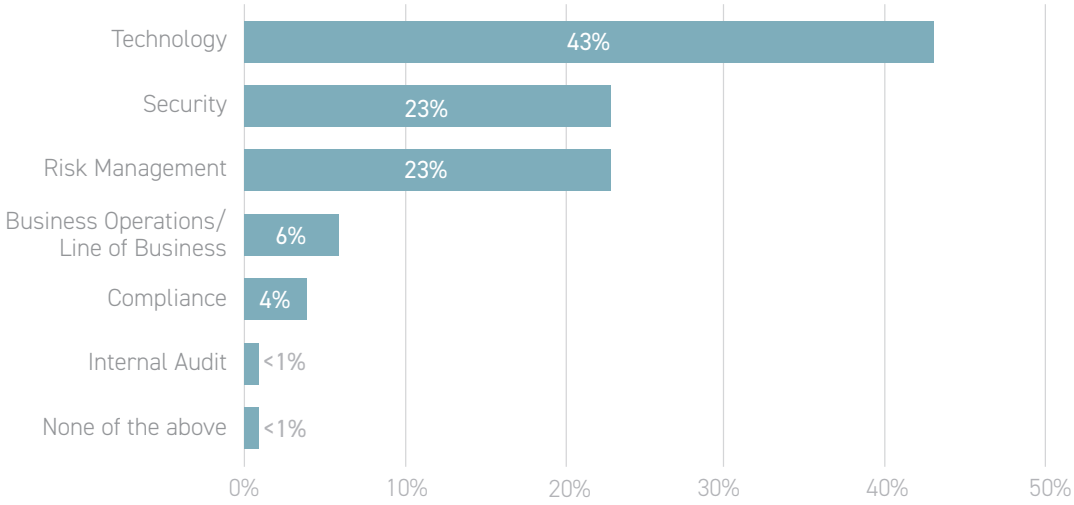

Source: AuditBoard

MOVING TO THE NEXT LEVEL WITH AGENTIC AI

Agentic AI is the next evolution of AI. Think of Gen AI as the musical instruments in an orchestra; each instrument—strings, woodwinds, and so forth—plays a specified part of the whole. Agentic AI, on the other hand, is like the conductor of an orchestra. It receives its instructions (the musical score) and then directs the various components (instruments) to create the result. Some current examples of Agentic AI include:

MULTIMEDIA CREATION

Although Gen AI can produce text, images, and video, agentic AI can aggregate and delegate subtasks like research, text generation, image selection and design to other AI systems. Agentic AI can process immense datasets—including transactional records, demographic shifts, economic indicators, and even social media sentiment—to identify hidden patterns and forecast market trends. Incorporating these dynamic data streams allows the creation of interactive marketing materials, such as video walkthroughs of properties and comparisons with other available properties of interest for both the CRE firm and client.

KNOWLEDGE RETRIEVAL

Agentic AI improves knowledge retrieval by accessing information and taking action based on insights. To illustrate, consider IT helpdesk operations. Whereas earlier generation helpdesk chatbots could answer specific, well-defined user questions, agentic AI goes deeper: analyzing issues, offering options, or asking clarifying questions to narrow down information. Agentic AI tools can extract key terms, clauses, and dates from unstructured lease agreements and legal documents, replacing weeks of manual effort with minutes of automated analysis.

RISK REDUCTION AND SECURITY

Agentic AI can aid in enterprise security operations and risk reduction efforts by orchestrating the components of those activities. For example, AI agents in a security operations center can proactively scan for new and emerging threats, investigate anomalies, and automatically take corrective action without human intervention. Similarly, in risk management, AI agents can search for unusual activity, investigate those patterns to determine if they're truly fraudulent, and automatically respond as needed.

UTILITIES

Agentic AI is already being used in the utilities industry. For instance, utilities are testing AI agents' ability to assess, triage and organize responses to disasters, such as hurricanes and wildfires. The agent can analyze data to rate infrastructure damage and its effect on individuals and communities; plan and schedule rescue and repair work; and route the workers and materials needed to complete repairs on time. This can dramatically accelerate recovery times, potentially saving lives in the process.[5] From predictive maintenance systems that anticipate equipment failures before they occur, to intelligent energy management systems optimizing consumption based on real-time occupancy and weather, Agentic AI is making buildings smarter, more efficient, and more sustainable.

Agentic AI is the next evolution of AI. Think of Gen AI as the musical instruments in an orchestra; each instrument—strings, woodwinds, and so forth—plays a specified part of the whole. Agentic AI, on the other hand, is like the conductor of an orchestra.

REASONS FOR CAUTION

Agentic AI systems pose regulatory, security, data, and workforce challenges—not unlike GenAI systems. These problems are arguably even more important and challenging due to the increased complexity of agentic AI systems.

ERRORS/HALLUCINATIONS

Hallucinations refer to instances where LLMs generate outputs that are factually incorrect, nonsensical, or otherwise not grounded in reality. Essentially, a hallucination occurs when the AI "makes things up" or produces information that is not supported by its training data or real-world evidence. This can manifest as inaccurate information, nonsensical text, false citations and/or references, or even impossible visual content.

ETHICAL DILEMMAS

Ethical dilemmas may include bias and discrimination, privacy violations, and accountability. As discussed earlier, AI systems learn from data, so if the data is biased, the AI decision making will be biased as well. Because Agentic AI uses multiple systems, each drawing from their own specific data pools, there is the potential for data collection that may violate privacy laws. Traditional privacy laws require collecting only necessary data, but AI often performs better with larger datasets. This creates tension between data minimization principles and AI effectiveness.

Finally, when an AI system makes an unethical or harmful decision, who is accountable? If, say, the AI system used Personally Identifiable Information (PII) as part of marketing outreach to potential buyers, would the management team of the CRE firm be held responsible for the data breach? As shown in *Exhibit 2*, this has been an ongoing issue for a long time.

Problems are arguably even more important and challenging due to the increased complexity of agentic AI systems.

EXHIBIT 2: IBM PRESENTATION, 1979

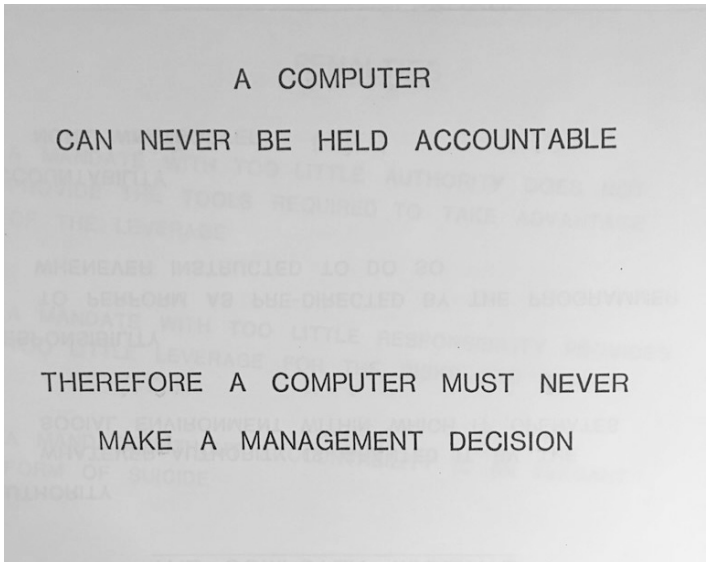

Source: Simon Willison

PRESSURE OF A VOLATILE LANDSCAPE[6]

REGULATIONS ARE MOVING FAST

The landscape of privacy and AI law continues to evolve at a breakneck speed. Regulations are evolving faster than ever, but are still running behind AI development and usage.[7] Even so, organizations such as The Center for AI Safety are still advocating for more regulation in the AI space.[8] For instance, there is no US equivalent to the European Union Artificial Intelligence Act. Instead, in the US, the federal government has stepped back from AI governance and ethics and states have started taking a more active role in defining the regulatory framework and enacting their own laws, as the following map illustrates. The result so far has been a piecemeal approach to AI regulation which can complicate compliance across multi-state portfolios.

Regulations are evolving faster than ever, but are still running behind AI development and usage.

EXHIBIT 3: CURRENT STATE OF AI LEGISLATION IN THE US

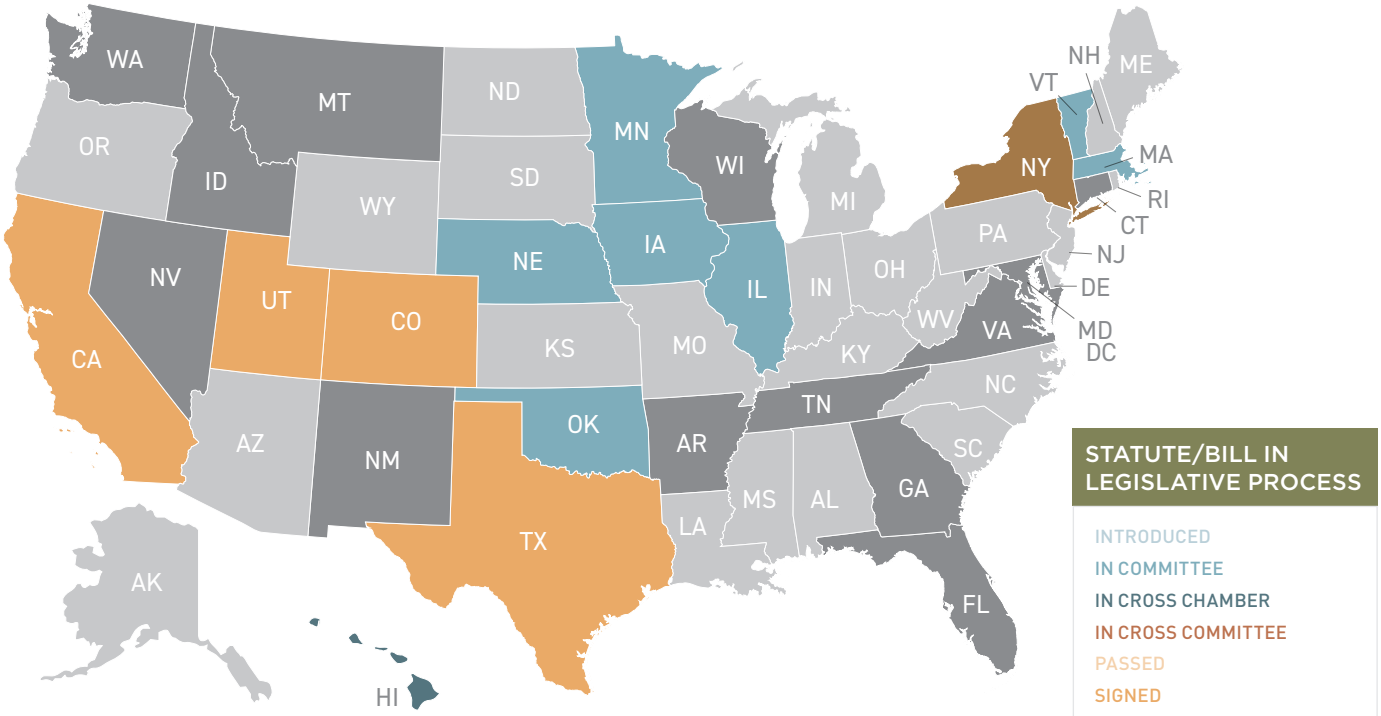

Source: IAPP

Last updated 1 July, 2025

NEW REGULATIONS ARE COVERING MULTIPLE DOMAINS

Between the marketing appeal of various building labels (e.g., LEED, BOMA, WELL, WiredScore, BREEAM, etc.) and investor demands for more ESG disclosures, landlords and managers collect a lot of data at the building level without full control of the data lifecycle. AI makes ESG reporting economically viable even when not legally required, allowing companies to build capabilities gradually and position themselves advantageously as requirements inevitably expand. At this point, these demands are more market than regulation driven but represent an altogether new pressure point where AI can help and where rail guards need to be established.

The intersection between data privacy and AI creates complex compliance challenges as AI systems fundamentally change how personal data is processed, analyzed, and used for decision-making, thus amplifying the need for transparency and disclosure from market players.

The intersection between data privacy and AI creates complex compliance challenges as AI systems fundamentally change how personal data is processed, analyzed, and used for decision-making, thus amplifying the need for transparency and disclosure from market players.

For instance, if asked by customers, tenants, investors or regulators, CRE professionals need to disclose what AI models are used, and what data sources are fed into AI systems. AI can derive unexpected insights from data that was originally collected by organizations for other, non-AI purposes. For example, if a real estate organization collected property data for valuation purposes, and shared that data with an AI that subsequently discovered potential patterns for predicting tenant creditworthiness, this unintended use and application of the data may constitute a new purpose requiring the organization to request additional consent from its stakeholders and customers, or related disclosures.

AI CAN BE CONTROLLED

AI risks span across an entire organization. Thus, like with other enterprise-wide risks, real estate companies need to establish a company-wide accountability system. This system should have tiered roles, and a clear governance hierarchy from operational implementation at the bottom to board oversight at the top. Broadly, this implementation would involve:

1. Assigning clear roles for AI model development, validation, monitoring, and approval; develop data processing agreements with AI vendors and ensure they meet the corporate privacy procedures.

2. Developing policies for AI model risk management including acceptable use cases, performance thresholds, escalation procedures and implement regular bias testing

3. Creating an AI steering committee including risk management, compliance, IT, and legal representatives. The main difference here is that ethics of the AI agent needs to be monitored

EXHIBIT 4: AI GOVERNANCE STRUCTURE

Source: Authors

UNDERSTAND THE AI LIFECYCLE

Just like there is a data lifecycle, there is an AI lifecycle. It refers to the complete journey of an AI model from initial conception to final retirement. The lifecycle approach ensures a maintained control and compliance throughout the AI model's operational life, rather than just at the initial deployment. This is especially important in real estate which involves complex, long-term property asset management cycles and holding periods. It can be broken down in several phases, each one with its specific governance requirements:

- **Planning:** Defining the core business challenge to be solved by AI, including data "preparation," collection, checking/ cleaning, and labeling. This phase resembles the "old-fashioned- market research and feasibility studies." For instance, if the goal is to optimize rent levels, managers would start by collecting and verifying data such as historical rent rolls, market comparables, tenant move-out reasons, local employment rates, competitor pricing, property amenities, maintenance costs, and so forth.
- **Development:** In this phase, managers would evaluate platforms and decide to either select one or build a proprietary system. If building a custom solution, the model would be fed and trained using the data prepared in the planning phase.
- **Validation:** Testing the model's accuracy, fairness, and reliability. At this point, are the results acceptable and do they make sense? For instance, does the algorithm show a bias toward Class A properties while the portfolio is mostly Class B? If a submarket is particularly volatile or some properties have unique elements, human judgement is preferable, and flags or escalation procedures should be put in place.
- **Deployment:** Rolling the model into production where it will make actual business decisions. This includes integrating with existing systems, training staff on how to use AI outputs, including awareness of the tool's capabilities and limitations, establishing override procedures, and setting up monitoring systems. In this phase, it is critical to continue training staff and maintain expertise to use the tools effectively and know how to step in should they become unavailable.
- **Retirement:** All AI models at some point become outdated and need replacement. This phase involves safely decommissioning the old model, migrating to new systems, preserving documentation for regulatory records, and conducting post-implementation reviews to learn lessons for future projects and knowledge transfer. Attention should be given to data lifecycle: AI systems can be processing client data and storing them longer than necessary, increasing breach risks. The market has not reached this point yet. In the above example, even without retiring the AI model altogether or switching to a newer software, data related to tenants who have not renewed or buildings which have been sold should be deleted. Use cases and query criteria would need to be updated if market conditions changed (i.e. new zoning changes, a major relocation in the submarket…).

EXHIBIT 5: LIFECYCLE OF AI

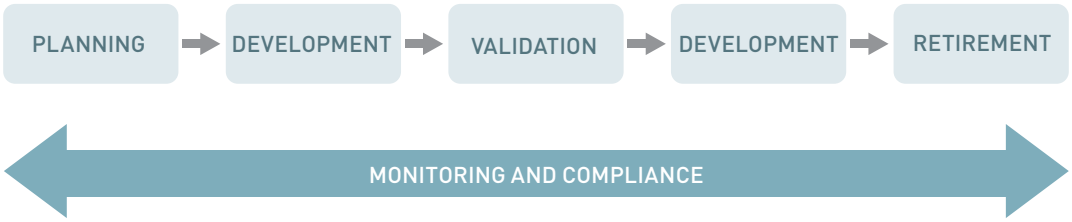

Source: Authors

ADAPTING THE CULTURE

Using AI tools requires a new corporate culture and a commitment to change. AI is not an "install and forget it" software. It requires significant planning and understanding of the processes and people that will be affected by its use.

Throughout the AI lifecycle, it is important to keep monitoring and compliance oversight. If using third-party AI tools, one should ensure vendors meet the corporate risk management standards and should require transparency into model methodology and regular performance reporting. Contingency plans for vendor service disruptions also need to be established, maintained, and included in the company's business continuity plan.

Finally, with all the capabilities these new AI tools can provide, it is critical that human involvement is included in any use. An AI tool should never be left to operate on its own. Ethical use of AI starts and ends with human oversight.

ABOUT THE AUTHORS

Marie-Noëlle Brisson, FRICS, and Michael Savoie, PhD, are Co-Founders of CyberReady, LLC, which provides cyber risk management, and state-of-the-art online and in-person training and assessments of the cyber risk profile of an organization's physical, behavioral, technical, and data privacy assets.

NOTES

[1] Now decides next: Generating a New Future. Deloitte's State of Generative AI in the Enterprise Quarter four report. January 2025; What is Artificial Intelligence? Accessed July 13, 2025. https://www.nasa.gov/what-is-artificial-intelligence/

[2] Finn, Teaganne and Amanda Downie. Agentic AI vs. Generative AI. Accessed July 13, 2025. https://www.ibm.com/think/topics/agentic-ai-vs-generative-ai

[3] https://www.britannica.com/technology/Google-Gemini

[4] https://go.auditboard.com/rs/961-ZQV-184/images/AB-EB-Conquering-Compliance-Navigating-the-Triple-Threat-of-a-Volatile-Regulatory-Landscape.pdf

[5] Pratt, Mary K. 10 real-world agentic AI examples and use cases. AI that doesn't just follow instructions but figures out how to get things done -- that's the promise of agentic AI, an emerging approach that's already changing some sectors. Published 07 March 2025. https://www.techtarget.com/searchenterpriseai/feature/Real-world-agentic-AI-examples-and-use-cases

[6] Conquering Compliance: Navigating the Triple Threat of a Volatile Regulatory Landscape. https://go.auditboard.com/rs/961-ZQV-184/images/AB-EB-Conquering-Compliance-Navigating-the-Triple-Threat-of-a-Volatile-Regulatory-Landscape.pdf. Accessed July 31, 2025.

[7] Brisson, Marie-Noelle, and Michael Savoie. Cyber Risk Vigilance. Summit Journal, Issue 17. 2025. AFIRE publishing. https://www.flipsnack.com/afire/summit17.html

[8] https://safe.ai/

AI is not an “install and forget it” software. It requires significant planning and understanding of the processes and people that will be affected by its use.

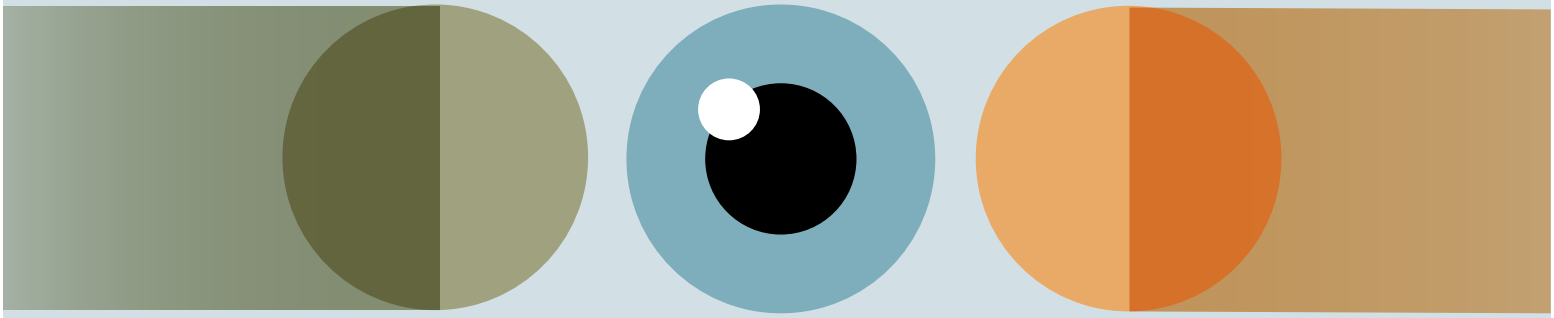

Ethical use of AI starts and ends with human oversight.