RESEARCH SERIES · VOLUME 1 2026

DECISION SCIENCE SOLUTIONS

AI FOUNDATIONS

What Every Actuary Should Know

Part 1 - A conceptual framework and common language for the global actuarial profession

Authored by

Asif John FIoA, ASAA, MSc
Chair, IFoA AI & Emerging Technologies Practice Board · Partner, ARGenesis

Alex Waite FIA, C.Act, CERA
Council Member & President-Elect of the IFoA · Council Representative on the AIET Practice Board

VERSION 1.00 · JUNE 2026

AUTHORS' PREFACE

From the Chair of the AIET Practice Board:

As the actuarial profession enters an era shaped by Artificial Intelligence (AI) and accelerated decision science, clarity has never been more important. Actuaries increasingly work alongside technologies that influence many aspects of our work, yet many members remain uncertain about what these systems are and how to assess them. This paper is a step towards establishing a shared foundation across the profession.

Our aim is to provide accessible explanations of modern AI, demystify key concepts and help members engage confidently with the tools, risks and governance expectations that now arise in practice. As Chair of the AIET Practice Board (PB), I believe this clarity supports stronger professional judgement, clearer communication with stakeholders and more responsible adoption of AI across insurance and financial services.

This paper is intentionally foundational. It establishes a shared conceptual framework and common language for discussing AI within actuarial work. More detailed practical guidance, including governance checklists, evaluation approaches and case studies is expected to follow in subsequent publications. This staged approach reflects the Practice Board's intention to support members progressively as AI use across the profession develops.

This paper also marks the starting point for a broader programme of education, research and horizon scanning that the AIET PB intends to deliver over the coming years. I hope it serves as a useful reference for members as we navigate a rapidly evolving landscape together.

Asif John FIoA, ASAA, MSc
Chair, IFoA AI & Emerging Technologies Practice Board
Partner, ARGenesis

From the IFoA Council representative on the AIET PB:

Actuarial science has always rested on a shared foundation: probability, modelling, professional judgement, clear communication, independent challenge and accountability. As new technologies reshape how decisions are made, the profession must be explicit about how that foundation extends to modern, data-driven systems. From an IFoA perspective, this work speaks directly to the future vision of the profession. Our public value has always been grounded in our ability to combine technical understanding with ethical judgement, proportionality and transparency, and challenge decisions where appropriate. There has never been a greater need for these skills than in the current era of technology.

Alex Waite FIA C.Act CERA
Council Member & President-Elect of the IFoA
Council representative on the AIET Practice Board

ACKNOWLEDGEMENTS

The authors would like to thank **Dhruva Divate, PhD** and **Mike Fenton, FIA**, for generously giving their time to review this paper. Their thoughtful comments, professional insight and constructive feedback helped strengthen the clarity, structure and technical quality of the work.

VERSION

This is version 1.00 of this document. We would like to encourage you to provide feedback to enable us to amend the document for future iterations. The authors can be contacted by emailing alex.waite@alexwaite.org or asif.john@argenesis.com

Terminology

The world of Artificial Intelligence (AI) is littered with domain-specific terminology; sadly, much of which is ambiguous. Some believe that this is deliberate so as to exclude *"outsiders"* from the AI domain. Throughout this paper, we have strived to explain what we mean when using various technical terms (which we believe to be aligned to the approach endorsed by the actuarial profession). Generally, capitalised terms such as Validation and Hallucination are used in the specific sense which is defined in the Glossary. However, there are terms that remain ambiguous and cases where general words can be used in more than one sense – this is somewhat inevitable!

COPYRIGHT

Contents

1 Executive Summary

Artificial Intelligence (AI) is increasingly being discussed and adopted across insurance, pensions and wider financial services, including within actuarial work. However, the breadth of AI-related terminology and the pace of technological change have made it difficult for many practitioners to develop a clear, shared understanding of what different AI technologies are, how they differ and what their implications are for actuarial practice.

This paper provides a non-technical foundation for actuaries seeking clarity on modern AI concepts and their relevance to actuarial work. It is intended for a broad professional audience and does not assume prior technical knowledge. The paper does not provide implementation guidance or endorse specific tools. Instead, it focuses on building conceptual understanding to support informed professional judgement, governance and oversight.

The paper introduces the contemporary AI landscape, as it stands in 2026. We introduce key technologies such as Machine Learning, Deep Learning, Large Language Models, Generative AI, Retrieval-Augmented Generation (RAG) and Agentic systems in plain language. The paper also presents a simplified AI lifecycle aligned with established actuarial approaches to model development, Validation (as defined in the glossary), deployment and ongoing monitoring based around the traditional actuarial control cycle.

Practical applications of AI across the work of actuaries are illustrated through both general business examples and specific case studies focussed on the specific specialisms in which we work as a profession. The paper also highlights key risks and limitations associated with AI, including data quality challenges, bias, bad actors, model performance degradation over time and the risk of incorrect or misleading outputs from language-based models.

Critically, the paper emphasises the importance of strong governance, clear accountability and ethical practice when AI is embedded in decision-making processes. It outlines priority areas for actuaries to focus on, helping the profession engage with AI in a way that is responsible, proportionate and consistent with actuarial standards.

Finally, the paper turns its gaze towards the horizon as we look at what emerging technologies might impact the future work of actuarial profession over the longer term.

Asif John & Alex Waite
17 May 2026

2 Introduction — Why the Profession Needs Clarity Now

Artificial Intelligence (AI) is becoming increasingly visible in actuaries' work within risk governance and all of the specialist domains of actuarial work. This includes insurance (pricing, reserving, underwriting, claims management, capital modelling), pensions (governance, administration and member communications) as well as wider fields (healthcare, medical analytics and financial services).

While many of the underlying analytical techniques have existed for decades, recent developments, particularly in data availability, computing power and language-based models, have accelerated adoption and broadened the range of potential use cases. As a result, actuaries are increasingly expected to engage with AI-enabled systems, either directly or through oversight, often within regulated and high-stakes decision-making environments. At the same time, the rapid expansion of AI-related terminology (often used inconsistently) has made it more difficult for practitioners to distinguish between different technologies, understand their limitations and apply appropriate professional judgement.

Against this backdrop, the purpose of this paper is to support actuaries in developing a clear, shared and non-technical understanding of modern AI technologies and their relevance to actuarial practice. It does not aim to provide technical implementation guidance or to endorse specific tools. Rather, it seeks to clarify key concepts, distinguish between different AI capabilities and provide a framework for informed engagement, challenge and oversight.

By grounding the discussion in real-world examples from across the profession and established actuarial principles, the paper aims to help practitioners assess where AI can appropriately support actuarial work, while recognising its limitations and the continued importance of professional judgement.

3 The AI Landscape — A Simple Overview

AI is best viewed as a layered set of techniques that can either be used on their own or combined together within a single workflow. Developing these layered techniques often takes considerable time and fully understanding the overall model can often be assisted by understanding the development journey, or “lifecycle” of the AI system. In this section we introduce the overview of how all these systems might fit together and then in Section 4 we will look at each part of the process in greater detail and expand upon the related legal, regulatory and compliance considerations.

3.1 The AI Pipeline — A Simple Lifecycle View

Most AI systems follow a familiar lifecycle: define the problem, collect and understand data, prepare the data, develop/train the model, validate performance and risks, deploy into processes and monitor over time, with outcomes and emerging evidence feeding back into reassessment of purpose, assumptions and model suitability. Risks typically arise from unclear objectives, weak data, poor Validation or changes in underlying patterns (Data Drift). The core discipline of the actuarial control cycle remains essential, including clear purpose, strong data, appropriate and documented assumptions, model Validation and ongoing oversight to ensure timely adjustment in the light of emerging evidence.

In Section 6 of this paper, we look in detail at the AI Pipeline.

3.2 The AI Stack of Technologies — A Practical View

A practical way to view modern sophisticated AI systems is as a stack: data/infrastructure → predictive modelling → application layers → output and orchestration. Higher layers depend on the quality and controls of the layers beneath, with governance, human oversight and feedback loops applying across all layers of the stack.

Taking this step by step (see figure 1 overleaf): at the base sit data and infrastructure, which determine what is feasible and how reliable outputs can be. On top of this are Machine Learning methods that learn patterns from historical data, with Deep Learning enabling stronger performance on complex and unstructured data. Large Language Models (LLMs) are a Deep Learning approach focused on language tasks, while Generative AI (GenAI) more broadly covers models that produce new content (often text or images in business settings). The reliability of GenAI in organisations is commonly improved through Retrieval-Augmented Generation (RAG), which grounds the AI output in approved sources. At the upper end, Agentic AI systems combine models, retrieval and tools to complete multi-step real-world tasks under defined controls.

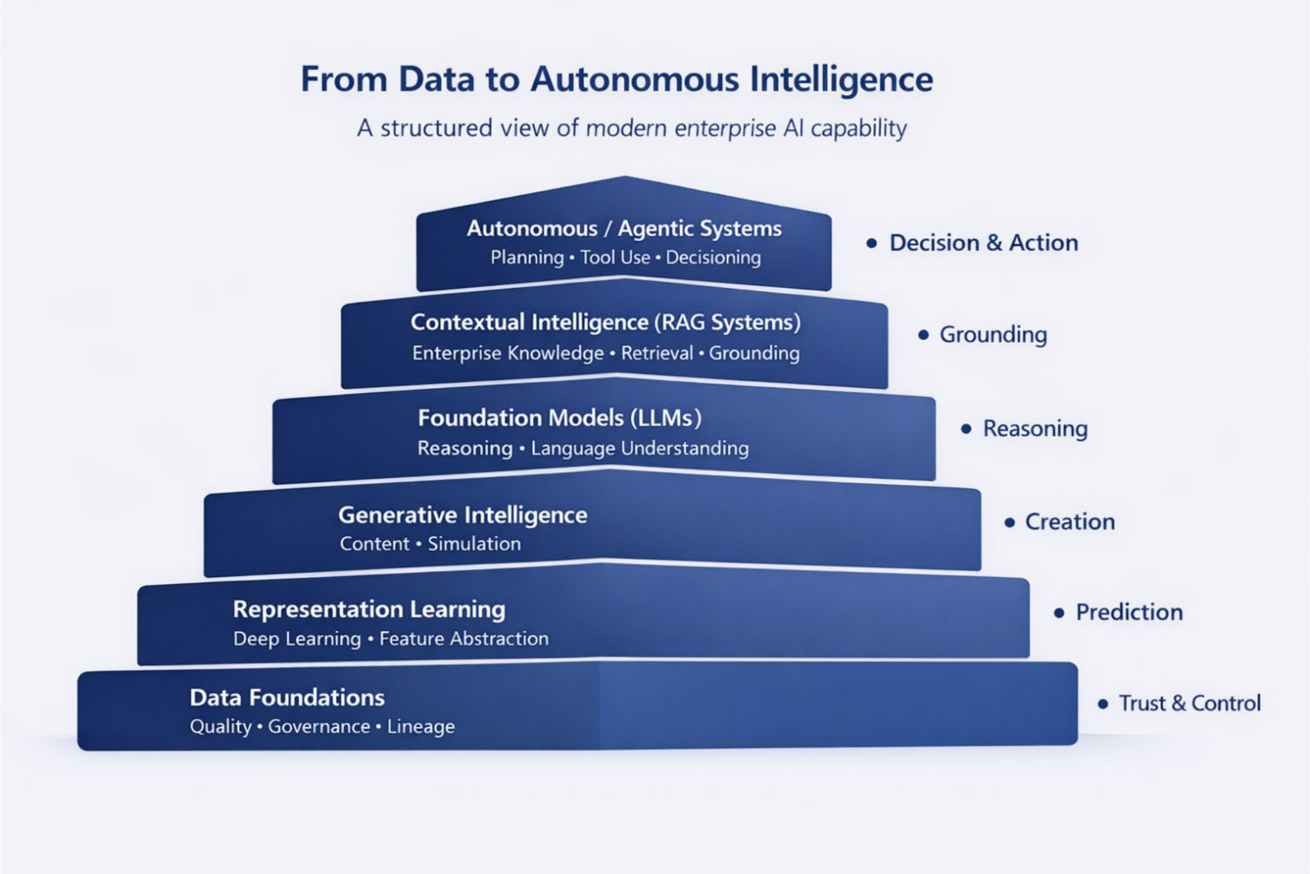

Figure 1 above provides a high-level illustration of this layered AI stack. More detailed explanations of each component are provided in Section 4.

4 Core AI Technologies Explained for Actuaries

This section summarises the main AI technologies typically encountered in actuarial contexts. It focuses on what each technology is designed to do, where it is typically applied and the key limitations to consider in practice, rather than how it is implemented.

4.1 Machine Learning (ML)

Machine Learning refers to a broad class of techniques that use data to identify patterns and relationships, which can then be applied to make predictions, classifications or rankings. Many Machine Learning approaches build on concepts that are already familiar to actuaries, such as regression modelling and experience analysis, but allow for greater flexibility in handling non-linear effects, interactions and large numbers of input variables. Whether ML is included within the scope of AI is a debatable point, but we include it within the definition for completeness.

Example Use Case: Netflix

Machine Learning is often used to make statistically based *"predictions"*; for example, for decades Netflix have used ML to predict what movies and TV shows viewers might like, based on their previous viewing habits. The algorithm might show that people who binge-watched *Bridgerton* also highly rated *The Crown*; therefore for a viewer who has already watched *Bridgerton*, it is worth suggesting *"You might also like The Crown"* as this will enhance the user experience (and hence increase retention and ultimately profitability).

The mathematics required for this use case is covered within the current actuarial syllabus.

Generalised Linear Models (GLMs) can be viewed as a form of supervised learning, where historical data with known outcomes are used to estimate relationships between risk factors and responses. Experience rating similarly involves learning patterns from past data to inform future expectations, while portfolio segmentation and risk grouping draw on ideas closely related to clustering and pattern recognition. Framing Machine Learning in this way helps demystify the techniques and reinforces that, in many cases, they represent extensions of established actuarial thinking rather than a departure from it.

In actuarial practice, Machine Learning techniques are commonly used to support pricing models, portfolio segmentation, lapse or churn analysis, fraud detection and risk scoring. Methods such as decision trees, gradient boosting and regularised regression can be particularly useful where relationships between variables are complex or where traditional model structures become difficult to specify explicitly.

While Machine Learning models can improve predictive performance, they may also be less transparent than simpler approaches. As a result, careful Validation, documentation, monitoring and clear communication of limitations remain essential, particularly in regulated decision-making contexts.

Common ML methods that are regularly found in actuarial work include:

- Decision trees and random forests, which split data into groups based on key characteristics and can support classification or prediction tasks.
- Gradient boosting methods, which combine many simple models sequentially to improve predictive accuracy.
- Regularised regression, which adds constraints to regression models to reduce overfitting and improve stability.
- Clustering algorithms, which group similar records together without needing predefined labels.
- GLMs and other regression-based methods, which model relationships between explanatory variables and outcomes.

Whilst all areas of actuarial work may benefit from this technology, ML is particularly common in insurance work, in particular:

- Predicting claim frequency and severity for pricing
- Identifying policies at high risk of cancellation or non-renewal
- Segmenting portfolios into more homogeneous risk groups
- Detecting potential large loss claims early based on claim characteristics

Benefits:

- Improved predictive accuracy compared with simple linear models
- Ability to incorporate more complex interactions and non-linear relationships
- Useful for high-dimensional data (many rating factors or predictors)

Limitations and cautions:

- Models can be harder to interpret than traditional actuarial methods
- Risk of overfitting if Validation and governance are weak, meaning the model may learn noise or unusual features in the training data rather than genuine patterns, leading to poor performance on new data
- Dependence on data quality and representativeness

4.2 Deep Learning (DL)

Deep Learning is a subset of Machine Learning that uses multi-layer neural networks (A neural network is a computational structure, loosely inspired by biological neurons, in which data passes through layers of interconnected processing units to identify patterns or make predictions.) to model highly complex relationships within data. These techniques are particularly effective when dealing with unstructured data, such as text, images or audio or with high-dimensional structured data involving complex interactions between variables.

Deep learning methods are increasingly applied to tasks such as processing scanned documents, assessing insurance damage from images or extracting information from large volumes of free text. These applications can significantly reduce manual effort and enable new forms of analysis that were previously impractical. However, Deep Learning models can be computationally intensive and more difficult to interpret than traditional actuarial models. Their use therefore requires careful consideration of proportionality, Explainability and governance, particularly when outputs influence material decisions.

Business examples relevant to actuaries might include:

- Processing scanned documents (e.g. legal documentation, medical reports, engineering surveys) using Optical Character Recognition (OCR) and text interpretation
- Analysing feedback reports or complaint logs from individuals in order to identify emerging risks or potential fraud
- Assessing damage from photos in property or motor claims
- Classifying and routing incoming pension scheme member emails

Benefits:

- Excellent performance on tasks where traditional models are weak
- Ability to extract features automatically without manual *"feature engineering"*, meaning the process of manually selecting, creating or transforming input variables for use in a model

Limitations and cautions:

- Generally, less interpretable than simpler models
- Computationally intensive and data hungry
- Vulnerable to subtle Data Drift and changes in underlying behaviour

Example Use Case: YouTube

A widely recognised example of Deep Learning is content recommendation on platforms such as YouTube or Spotify (and indeed Netflix in the modern era). These systems do not just rely on simple rules or past selections but instead learn complex patterns in user behaviour over time, including the sequence and context of interactions and even characteristics of the content itself (such as audio or video features).

Similar approaches can be applied in actuarial contexts, for example to better understand policyholder behaviour, predict lapse or retention and tailor products or engagement strategies based on observed patterns.

4.3 Large Language Models (LLMs)

Large Language Models (LLMs) are a class of Deep Learning models trained on large volumes of textual data to learn statistical patterns and structures in language. This enables them to generate text, summarise documents, classify information and respond to queries in natural language.

In actuarial and general business contexts, LLMs are commonly used to assist with tasks such as summarising reports, extracting information from documentation, preparing meeting summaries/actions, drafting narrative explanations or supporting search and retrieval across large document collections. They can act as a flexible interface between users and complex information sources.

It is important to recognise that LLM outputs are generated based on learned patterns rather than verified facts. As a result, they may produce incorrect or misleading responses if not appropriately constrained. Risks such as hallucinated content or vulnerability to prompt manipulation (where users deliberately craft inputs that cause a model to bypass intended controls, ignore instructions or produce inappropriate outputs) highlight the need for controls, review processes and supporting governance when LLMs are used in professional environments.

Actuarial-relevant examples:

- Summarising complex claim files into concise narratives for review
- Converting a pension Trust Deed & Rules into a member benefit specification
- Extracting key fields from broker submissions or proposal forms
- Turning technical model documentation into plain-language explanations for non-specialists

Benefits:

- Dramatically reduces manual effort in reading and summarising text

- Enables “chat-style” interaction with documentation or processes
- Makes specialised material more accessible to wider stakeholders

Limitations and cautions:

- LLMs can generate incorrect but highly plausible responses (*“Hallucinations”*)
- Must be used with strong guardrails, particularly in regulated contexts
- Needs to be coupled with organisational knowledge (e.g. via Retrieval-Augmented Generation, described in Section 4.5) for context-specific answers and improved accuracy

Example Use Case: ChatGPT

A widely recognised example of Large Language Models is conversational AI tools such as ChatGPT (or Claude, Gemini or others). These systems allow users to interact with large volumes of information through natural language, for example by asking questions, requesting summaries of documents, or generating draft responses. Rather than retrieving pre-written answers, the model generates responses based on patterns learned from extensive text data, enabling flexible and context-aware interaction.

Similar capabilities can be applied in actuarial contexts, where LLMs can support tasks such as summarising reports, explaining technical results in plain language, drafting documentation, or enabling users to query large collections of actuarial or business information through a conversational interface.

Key legal, compliance and third-party risk considerations in relation to LLMs are set out in Section 5.

4.4 Generative AI (GenAI)

Generative AI refers to models designed to produce new content such as text, code or structured outputs based on patterns learned from data. In practice, GenAI applications used in many of the business contexts familiar to actuaries will be built on LLMs, particularly for text-based tasks, including coding.

For actuaries, GenAI can be useful in drafting or reviewing documentation, producing preliminary analyses, summarising model outputs or translating technical material into clearer language for non-specialist audiences. These capabilities can improve efficiency and consistency, particularly for repetitive or standardised tasks. However, generative outputs should not be treated as authoritative without review. Professional judgement remains essential to ensure that outputs are accurate, appropriate and aligned with regulatory and ethical expectations.

Relevant examples for actuaries include:

- Drafting *"movement commentary"* based on key drivers (whether this be changes in triangles for general insurance or an analysis of surplus arising for pensions)
- Generating first drafts of documentation, meeting notes, or model change reports
- Creating skeletons for insurance policy wording comparisons or endorsement summaries
- Creating a first draft of a letter to a pension scheme member or a member video using an AI avatar
- Assisting with code snippets for data processing, reporting or model implementation (R, Python, VBA, SQL, or other coding languages)

Benefits:

- Reduces time on repetitive drafting tasks
- Provides consistent initial templates that can be refined by humans
- Encourages better documentation through lower effort barriers

Limitations and cautions:

- Outputs still require professional review and sign-off
- Beware confidentiality policies and constraints for using LLMs for confidential work
- Risk of subtle errors in technical or regulatory content (and notably a different class of error than one might typically expect from a human doing the same work)
- Potential for inconsistent style if not carefully guided

Example Use Case: Drafting with an LLM

A widely recognised example of Generative AI is tools such as ChatGPT, or an in-house equivalent LLM model, which can generate a first draft of an email, report or explanation from a short prompt or a few bullet points. Similarly, tools such as GitHub Copilot (or Claude Code) can generate code from natural language instructions or partial inputs. These examples illustrate how Generative AI can create new content rather than simply classify or predict.

In actuarial contexts, similar capabilities can be used to draft movement commentary, prepare first versions of documentation or communications and assist with code for data processing, reporting or model implementation, with appropriate human review and oversight.

Key legal, compliance and third-party risk considerations in relation to GenAI are set out in Section 5.

4.5 Retrieval-Augmented Generation

Retrieval-Augmented Generation (RAG) combines GenAI with information retrieval from defined data sources or document repositories. Before producing an output, the system retrieves relevant material from an approved knowledge base and uses it to ground the response.

In actuarial contexts, RAG can support tasks such as querying pension scheme documentation, insurance policy documentation, underwriting/reserving guidelines or other governance materials using natural language. By anchoring outputs to authoritative internal sources, RAG approaches help reduce the risk of unsupported or fabricated content.

Effective use of RAG depends on the quality, scope and governance of the underlying document sets. Maintaining accurate and up-to-date source material is therefore a key control.

This reduces Hallucinations and explicitly ties responses to internal, authoritative sources.

Actuarial-relevant examples:

- Answering underwriters' questions based on internal pricing guidelines and manuals
- Providing consistent interpretations of policy wording for claims or coverage queries
- Answering queries from pension scheme members, via a chatbot, using references to the scheme documents and pre-existing member communications
- Acting as a natural-language interface for reserving methodology, assumptions and governance documents
- Supporting professional teams in quickly querying regulatory source material and proprietary explanatory notes

Benefits:

- Higher accuracy and traceability than *"pure"* LLM responses
- Easier to maintain alignment with internal standards and regulatory expectations
- Scales institutional knowledge beyond small expert groups

Limitations and cautions:

- Whilst reducing Hallucinations, RAG still suffers from the limitations and cautions of LLMs more generally. In particular, outputs still require professional review and sign-off.
- Requires well-curated and maintained document repositories.
- Governance is needed around which sources are indexed and how often they are updated.
- Where possible, responses should highlight which sources were used and the specific sections on the source that have been relied upon.

Example Use Case: Microsoft Copilot

A widely recognised example of Retrieval-Augmented Generation is enterprise AI assistants such as Microsoft Copilot, which allow users to query internal documents using natural language. For example, a user may ask a question about company policy or request a summary of a document, and the system first retrieves relevant information from approved sources before generating a response grounded in that material.

Similar approaches can be applied in actuarial contexts, where RAG can be used to query policy wording, reserving methodologies or governance documentation, ensuring that responses are based on authoritative internal sources rather than unsupported model outputs.

Key legal, compliance and third-party risk considerations in relation to RAG are set out in Section 5.

4.6 Agentic AI

Agentic AI systems extend beyond simple question-and-answer interactions by combining multiple components, such as language models, retrieval mechanisms and software tools to perform sequences of real-world actions toward a defined objective. These systems may plan steps, call external tools or models and adapt their behaviour within predefined constraints.

In pensions administration, AI is increasingly being used to support the handling of incoming member communications. A Deep Learning system can categorise each query and route it efficiently to the appropriate process or team; an Agentic AI system can potentially go further and initiate the administration process, seeking out the member's record and populating various fields in the database (provided appropriate oversight is maintained).

In insurance settings, agentic approaches may be used to support multi-step workflows, such as assembling data for pricing, running standard analyses or preparing draft reports. While such systems can improve efficiency, they also introduce additional complexity in terms of oversight and accountability.

Clear boundaries, human review points and robust governance are essential when Agentic systems are used in actuarial processes.

Relevant examples include:

- A life insurance underwriting agent that:
    1. Reads a broker submission
    2. Retrieves exposure and claims history
    3. Checks appetite and exclusions

4. Applies rating logic
5. Produces a draft quote and referral notes

- A pensions administrator agent that:
  1. Reviews an incoming member email and categorises it as *"change of address"*
  2. Identifies the member number and opens the appropriate member record
  3. Verifies the old and new address details and checks their Validity
  4. Flags to a human operator the required verification processes in this scheme
  5. Presents the administrator with a completed update request in the admin system, ready for confirmation and authorisation by the human administrator
- A house and contents insurance reserving agent that:
  1. Pulls claims and premium data
  2. Runs a set of standard actuarial methods
  3. Identifies unusual movements
  4. Drafts a commentary for the reserving committee pack
- A motor claims agent that:
  1. Takes the notification of loss input, potentially including photographs
  2. Classifies the claim, including a review of photographic evidence
  3. Checks coverage points
  4. Suggests a next action or triage pathway

Benefits:

- Automates end-to-end workflows, not just isolated steps
- Frees up actuarial, administrator, and underwriting time for higher-value judgement tasks
- Improves consistency and speed in routine processes

Limitations and cautions:

- Complex to design, test and govern
- Risk of automation-too-far without appropriate human controls
- Potentially increases opportunities for cyber-attacks due to reduced human oversight
- Clear accountability must be defined: who is responsible for outcomes?

Use Case: Autonomous AI assistants (e.g. "AI booking a trip for you")

A widely recognised example of Agentic AI is an autonomous digital assistant that can complete a multi-step task based on a user's objective. For instance, a user may ask the system to organise a trip and the AI will search for options, compare alternatives, make bookings and generate an itinerary, interacting with multiple systems along the way. Unlike traditional AI systems that respond to individual prompts, Agentic systems coordinate a sequence of actions to achieve a defined goal.

Similar approaches can be applied in actuarial and insurance contexts, where Agentic systems can support end-to-end workflows such as underwriting, claims handling or reserving processes, operating across multiple data sources and tools under appropriate human oversight.

Key legal, compliance and third-party risk considerations in relation to Agentic AI are set out in Section 5.

4.7 Synthetic Data

Synthetic Data refers to artificially generated data designed to resemble real datasets in their statistical properties, without representing actual individual records. It can be useful where real data are sparse, sensitive or fragmented, particularly in early-stage analysis or exploratory modelling. This can be particularly powerful in an AI context, which may require considerable training data, but Synthetic Data must also be used with care.

In actuarial work, Synthetic Data may support scenario testing, model development for new products, or analysis where confidentiality constraints limit access to real data. For example, in an insurance context, it may be used to simulate additional accident years for products with limited experience, generate alternative claim scenarios for stress-testing purposes or create de-identified datasets for internal training or external research collaboration.

In a pension context, Synthetic Data could be used to model the universe of UK occupational pension schemes to help assess the impact of a new risk transfer product before launching on the market.

Synthetic Data should not be treated as a substitute for genuine experience. While it can replicate broad statistical features of a dataset, it typically struggles to represent extreme events, rare combinations of risk factors or structural breaks all of which are critical considerations in actuarial analysis. As a result, Synthetic Data may understate extreme data points or fail to capture behaviours that emerge only under adverse or unusual conditions.

From an actuarial perspective, the use of Synthetic Data therefore requires careful Validation against real experience, clear documentation of limitations and appropriate professional judgement regarding its suitability for the intended purpose. Its role is best viewed as complementary to, rather than a replacement for, observed data.

Relevant examples include:

- Simulating additional accident years for new motor products with limited experience
- Simulating the universe of DB pension schemes based on available data points
- Generating alternative life insurance claim scenarios for capital and stress-testing purposes
- Producing de-identified datasets for internal training or external research collaboration

Benefits:

- Helps when data are sparse, sensitive or fragmented
- Supports scenario generation and robustness checks
- Can protect confidentiality while enabling analysis

Limitations and cautions:

- Poorly generated Synthetic Data may distort risk patterns
- Requires robust Validation against real data
- Should not be mistaken for genuine historical experience

Summary

These technologies are often combined within a single system, and commercial tools may obscure the distinction between them behind a unified interface. Understanding the underlying components helps actuaries assess where AI can add value, where risks arise and how established actuarial principles of Validation, documentation and professional judgement continue to apply.

Use Case: Self-driving car training (Waymo / Tesla)

A widely recognised example of Synthetic Data is its use alongside real-world data in training self-driving vehicles, where simulated driving environments are generated to supplement observed driving by exploring rare traffic scenarios, weather patterns and accident events. This allows models to be developed and tested on a wide range of situations, including those that are difficult, costly or unsafe to observe in real life.

Similar principles apply in actuarial contexts, where Synthetic Data can be used to simulate additional experience, explore alternative scenarios or support analysis where real data are limited or subject to confidentiality constraints.

5 Legal, Regulatory and Compliance Considerations

This section highlights common legal, regulatory and compliance issues that arise when AI is used in an actuarial context. It is intentionally practical and is not legal advice. The aim is to help actuaries identify where specialist input might be required and where controls should be strengthened before AI outputs are relied upon in business decisions.

5.1 Data protection, confidentiality and information security

AI systems often rely on large volumes of data, including personal data, commercially sensitive information and regulated records. Key considerations, which will be familiar to most actuaries from other data-handling activities, include:

- **Lawful basis and purpose limitation:** confirm the lawful basis for processing, the intended purpose(s) and whether secondary uses (for example, model training or vendor improvement) are permitted. A starting point for this will be obtaining clarity as to which legal system(s) apply, such as applicability of EU/UK General Data Protection Regulation (GDPR) or the EU AI Act.
- **Data minimisation and proportionality:** restrict data to what is necessary, particularly where decisions are material or regulated.
- **Special category and sensitive data:** health data, biometrics, criminal offence data or other sensitive categories may require enhanced safeguards and, in many cases, additional justification and controls.
- **Cross-border transfers and hosting:** understand where data is stored and processed, including vendor and sub-processor locations and ensure appropriate transfer mechanisms and contractual protections. In particular, unless specifically controlled there is no reason to assume that AI processing using cloud computing will be performed in the 'host' country and therefore additional data security tests may be required.
- **Confidentiality obligations:** ensure compliance with contractual confidentiality terms, professional obligations and internal policies (including restrictions on uploading information into third-party tools).
- **Security controls:** ensure people and systems only have the access they need, data is properly protected, critical duties are kept separate, activity is securely logged and security issues are detected and handled in line with the firm's information security standards.

5.2 Accountability, audit trail and record-keeping

Where AI supports decisions that impact customers, financial results or regulatory reporting, firms should be able to explain what happened, who approved it and what evidence supports it.

- **Clear ownership:** define business ownership, model ownership, user responsibilities and escalation routes (including when legal/compliance review is mandatory).
- **Decision traceability:** maintain an audit trail that links outputs to inputs, data versions, model versions, prompts, retrieved sources (for RAG) and human approvals.
- **Record retention:** ensure records are retained in line with regulatory and internal retention requirements and that outputs are reproducible where needed for audit or challenge.
- **Explainability and communication:** ensure outputs used in decision-making can be communicated to stakeholders in an appropriate form (including limitations and uncertainty).

5.3 Third-party, outsourcing and vendor risk

Many AI capabilities are delivered through third parties (model providers, cloud vendors, SaaS platforms, data providers). This creates additional compliance obligations.

- **Due diligence and contracts:** assess security, resilience, data handling, sub-processors, audit rights and termination/exit provisions.
- **Use of customer data:** confirm whether the vendor uses submitted data to train models or improve services; ensure opt-out, segregation and deletion mechanisms are understood and enforceable.
- **Processing location**: as noted above, unless specifically controlled there is no reason to assume that AI processing in the cloud will be performed in the 'host' country.
- **Change management:** vendors may update models frequently; controls are needed to detect material changes, reassess performance and re-approve where necessary.
- **Operational resilience:** assess service availability, incident reporting timelines, fallback procedures and the impact of outages on critical processes.

5.4 Conduct, fairness and discrimination risk

AI can amplify existing biases in data or create unfair outcomes through proxy variables and complex interactions. This is particularly sensitive in pricing, underwriting, claims and customer communications.

- **Outcome testing:** assess fairness using outcome-based monitoring (not just model performance metrics), including segmentation impacts and potential disparate outcomes.
- **Use-case sensitivity:** apply heightened scrutiny where decisions affect vulnerable customers, claims handling, complaints, cancellations, renewals or access to cover.
- **Human oversight and challenge:** ensure meaningful review is possible, including clear thresholds for human approval and exception handling.
- **Customer communications:** ensure explanations and disclosures are clear, accurate and consistent with conduct expectations.

5.5 Intellectual property, licensing and content risks

GenAI can create Intellectual Property (IP) and copyright issues, especially where prompts are used to guide outputs to resemble training examples or where proprietary content is used as input.

- **Input rights:** confirm the firm has rights to use any documents, datasets or third-party materials provided to AI tools (including internal client materials and licensed publications).
- **Output ownership and reuse:** clarify ownership, licensing and permitted reuse of generated outputs under vendor terms.
- **Risk of reproducing protected text:** implement controls to reduce verbatim reproduction and to flag high-risk uses (for example, drafting external-facing materials).
- **Inappropriate content:** some GenAI models have become notorious for producing content that is not appropriately controlled (for example, producing offensive language or images).
- **Confidential information leakage:** restrict sensitive documents to approved systems with appropriate access controls and logging.

5.6 Operational compliance and acceptable use

Practical controls are often the difference between safe experimentation and unmanaged risk.

- **Approved tools and permitted data:** define what tools may be used, for what purposes and what data types are prohibited or restricted.
- **User training:** ensure staff understand limitations (including Hallucination risk), safe prompt practices and when outputs must be verified.
- **Quality control:** require Validation checks proportionate to the materiality of the decision; define what "good enough" looks like for each use case.
- **Incident management:** establish processes for handling errors and 'near-misses' (for example, inappropriate outputs, security incidents, data leakage and customer impact) —aligned to existing risk and compliance reporting.

5.7 Mapping to regulatory expectations and internal governance

Firms should avoid treating AI as a special case outside established governance. Instead, AI should be mapped to existing frameworks for model risk, operational risk, data governance, outsourcing and conduct.

- **Align with model governance:** apply clear standards for Validation, documentation, approval and monitoring, scaled to materiality.
- **Align with risk appetite:** define *"no-go"* areas, approval thresholds and controls for high-impact decisions.

- **Regulatory engagement readiness:** ensure appropriate resources exist to support internal audit, external audit and regulatory review, including documentation of design choices and control effectiveness.

Closing principle: AI does not shift accountability. Even where systems are automated, responsibility for decisions, disclosures and outcomes remains with the firm and its professionals. Actuaries therefore have a key role in ensuring AI use is transparent, controlled and evidence based.

6 The AI Lifecycle — A Simple Pipeline View

Although AI systems can appear complex, most follow a broadly consistent lifecycle from conception to deployment and ongoing use. Understanding this lifecycle is particularly important for actuaries, as it aligns closely with established approaches to actuarial model development, Validation and governance. Viewing AI through a lifecycle lens helps clarify where risks arise, where professional judgement is required and where controls should be applied.

This lifecycle applies most directly where actuaries develop or train models themselves. GenAI relies on pre-trained third-party models and so the relevant development activity instead focuses on use-case definition, configuration (including prompts and retrieval mechanisms), controls and evaluation of outputs during inference, rather than model training.

The lifecycle typically begins with problem definition. At this stage, the purpose of the AI system is clearly articulated, including the decision or process it is intended to support. Success criteria, constraints and boundaries should be defined upfront, particularly in regulated environments. For actuaries, this mirrors the need to ensure that a model is appropriate for its intended use and that its outputs will be interpreted correctly.

The next stage is data collection and understanding. AI systems depend on relevant, representative data, which may include structured records such as policy or claims data, as well as unstructured information such as documents or text fields. Understanding data provenance, coverage, limitations and potential biases is critical. Weaknesses at this stage can materially affect downstream model performance and reliability.

This is followed by data preparation, which involves cleaning, validating and transforming data into a form suitable for modelling. Common activities include handling missing values, addressing inconsistencies and creating derived variables. For AI systems that use unstructured data, this stage may also involve text processing or document extraction. As with traditional actuarial work, careful data preparation is often a key determinant of model quality.

Once data are prepared, the model development and training stage begins. Appropriate techniques are selected based on the nature of the problem, ranging from traditional statistical methods to Machine Learning or Deep Learning approaches or combinations of these. Models are trained on historical data to capture patterns relevant to the defined objective. At this stage, design choices and assumptions should be documented clearly, as they affect interpretability and risk.

The evaluation and Validation stage assesses whether the model performs adequately for its intended purpose. This typically involves testing on data not used during training, examining performance metrics and checking robustness under alternative scenarios. In the context of AI, additional considerations may include sensitivity to input changes, stability over time and potential sources of bias. Performance degradation due to Data Drift where underlying patterns change over time is a particular risk that requires attention.

Following Validation, models may move to deployment and integration. This involves embedding the AI system into business processes, such as pricing workflows, claims triage or reporting pipelines. At this stage, it is important to define how outputs will be used, who is responsible for review and approval and how exceptions will be handled. Clear human oversight remains essential, particularly where AI outputs inform material decisions.

The next stage is monitoring and ongoing management. Once deployed, AI systems should be monitored to ensure they continue to perform as expected. This includes tracking model performance, identifying signs of Data Drift or misuse and determining when retraining or recalibration is required. Documentation and governance processes should be kept up to date to reflect changes in data, models or usage.

The final stage is continuously assessing the model for its appropriateness and predicting when it will no longer be *"fit for purpose"* and hence need to be retired. A key issue is that a replacement model needs to have been scoped, developed and tested in time to synchronise with the retirement of the old model. (Often the ideal process will be to parallel run both the old and the new models for a period of testing around the time of the switch over.)

Taken together, this lifecycle provides a structured way to think about AI systems in actuarial contexts. While the underlying techniques may differ from traditional models, the principles of clear objectives, sound data, Validation, documentation and ongoing oversight remain the same. Applying these principles consistently helps ensure that AI supports actuarial work in a responsible and reliable manner.

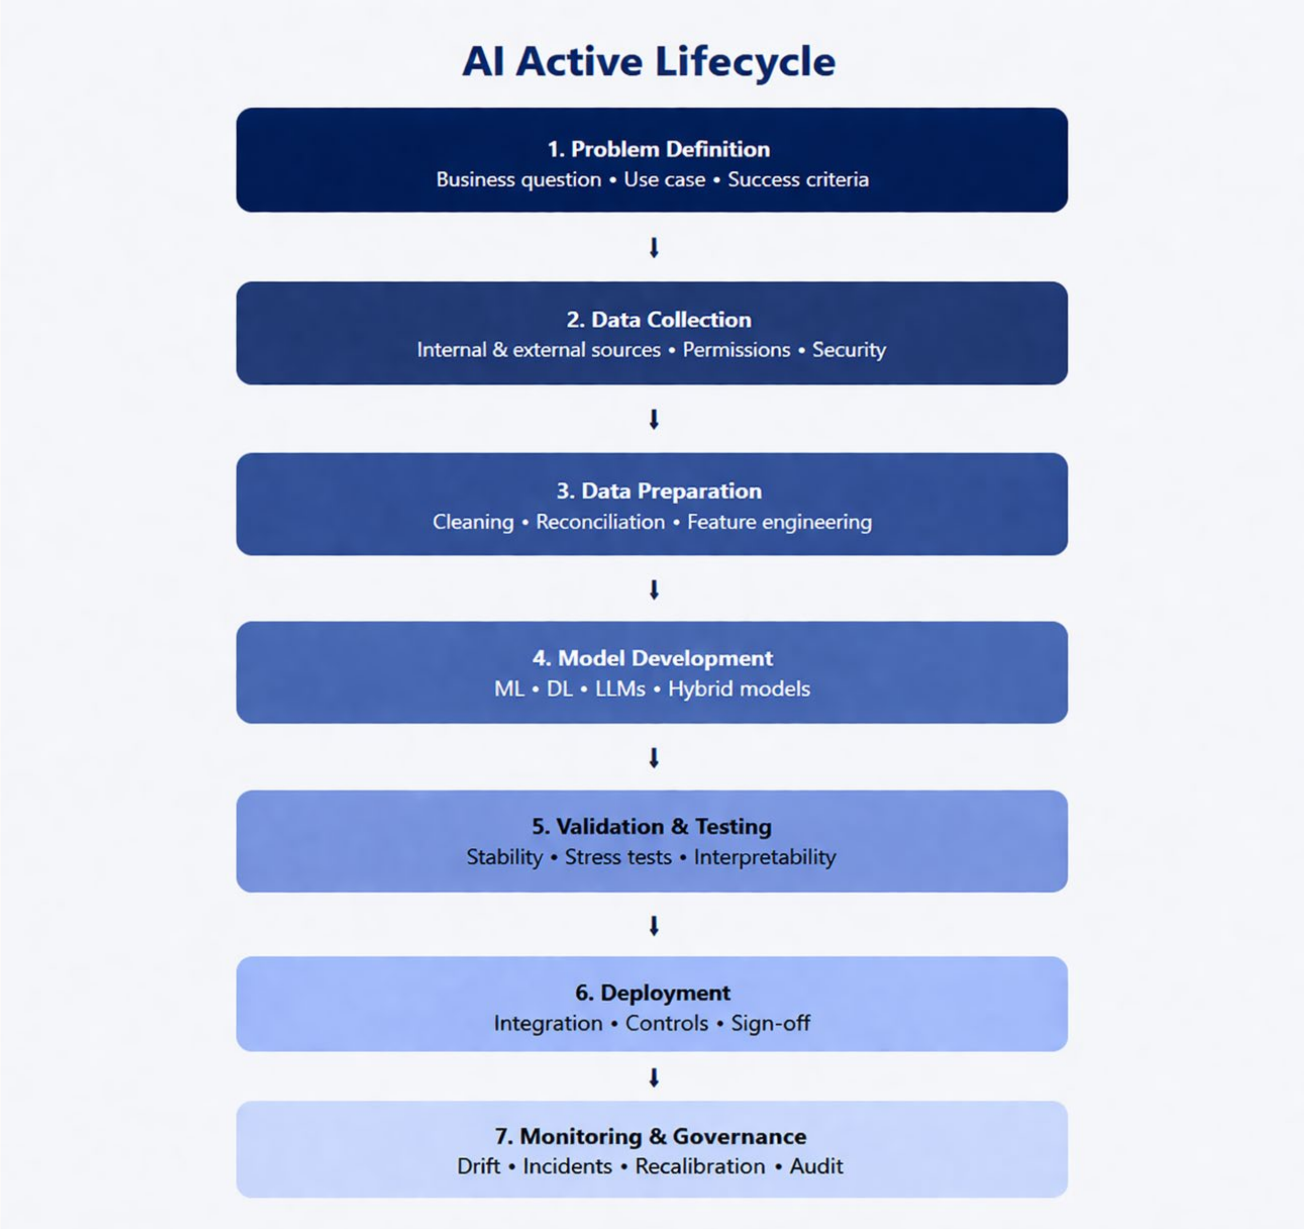

Figure 2 above provides a high-level illustration of a simple end-to-end cycle from problem definition → data → modelling → Validation → deployment → monitoring. It sets out the topics where actuarial judgement is most critical at each stage.

Actuaries can play a key role at each stage: framing the problem, interpreting outputs, shaping Validation, explaining behaviour and ensuring alignment with professional and regulatory standards.

7 Practical Applications Across Insurance

AI technologies are already being applied across many areas of actuarial work. The insurance examples in this section and the pension examples in Section 8 are intended to be illustrative rather than exhaustive, highlighting where AI is already commonly used in practice and how it can support actuarial judgement and decision-making. In all cases, the appropriateness of AI use depends on the context, materiality and governance arrangements surrounding the application.

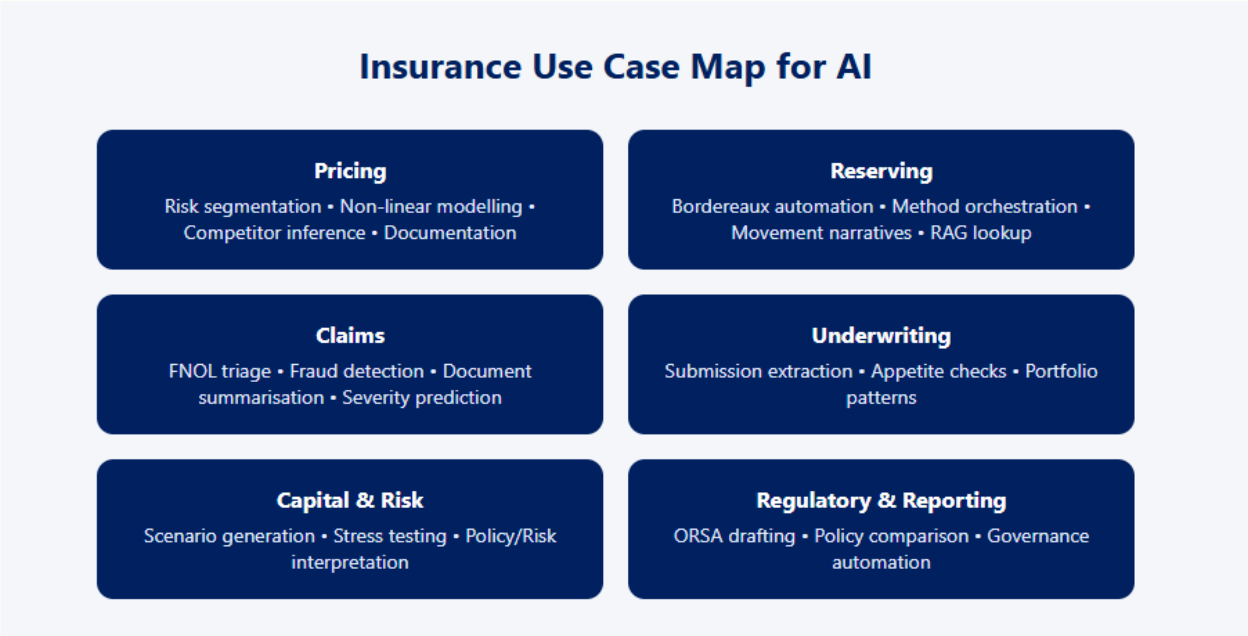

Figure 3 above provides a high-level grouped view of AI use cases across the insurance value chain. (Please note that some areas that are discussed separately in the sections that follow are combined in the figure for simplicity of presentation.)

7.1 Pricing

In pricing, AI is often used to support risk assessment, segmentation and rate optimisation. Machine Learning techniques can help identify complex relationships between rating factors and outcomes, particularly where interactions or non-linear effects are present. These methods may complement, rather than replace, more traditional actuarial approaches.

AI can also assist earlier in the pricing process by supporting data preparation and quality checks, such as identifying anomalies, inconsistencies or missing information in large datasets. In more advanced implementations, AI-enabled tools may help generate draft pricing documentation or summarise model changes for review, subject to actuarial oversight.

While AI can improve efficiency and predictive performance, pricing decisions remain subject to regulatory, ethical and commercial constraints. Transparency, fairness considerations and clear communication of assumptions continue to be essential.

- **Data preparation:**
  Using ML and DL to clean, standardise and enrich rating data, including free-text fields from broker notes or proposal forms.
- **Risk segmentation:**
  Applying ML models to identify homogeneous risk groups and non-linear interactions between rating factors.
- **Competitive analysis:**
  Estimating implied competitor rates from market quotes using ML techniques.
- **Scenario testing:**
  Using Agentic AI workflows to re-run pricing models under different assumptions and produce summary reports for pricing committees.
- **Documentation and explanation:**
  Using GenAI to draft pricing methodology descriptions and change logs.

7.2 Reserving

In reserving, AI is typically applied to support data handling, analysis and communication rather than to replace established actuarial methods. Machine Learning techniques may be used to explore patterns in claims development data, identify unusual movements or support segmentation of portfolios.

Language-based AI tools can assist with tasks such as summarising claims information, extracting insights from narrative data or drafting explanatory commentary for reserving reports. These applications can reduce manual effort and improve consistency, but they do not remove the need for actuarial judgement in selecting methods, setting assumptions or interpreting results.

AI-enabled reserving processes should be subject to the same governance standards as traditional approaches, including Validation, documentation and regular review.

- **Data ingestion:**
  LLMs and DL models can assist in reading bordereaux — detailed schedules of policies, premiums, claims or exposures shared between insurance parties — aligning fields and flagging anomalies or inconsistencies.
- **Method automation:**
  AI-enabled tools can orchestrate standard actuarial techniques (e.g. chain-ladder, Bornhuetter-Ferguson, stochastic methods), run them consistently and highlight unusual movements.

- **Narrative production:**
  GenAI can draft reserve movement narratives, explaining quarter-on-quarter changes, drivers and key risks, subject to actuarial review.
- **Knowledge retrieval:**
  RAG can support reserving teams by providing fast access to historical committee papers, model documentation and prior decisions for context.

7.3 Claims

Claims management is an area where AI adoption has been particularly visible. AI systems are often used to support triage, routing and prioritisation of claims, helping to direct cases to appropriate handling teams. Image and text analysis techniques may assist in assessing damage, extracting information from documents or identifying potentially fraudulent behaviour.

From an actuarial perspective, AI applications in claims can influence both operational outcomes and the quality of data used for downstream analysis. Actuaries therefore have an important role in understanding how AI-driven claims processes affect experience data, reserving assumptions and performance metrics.

As with other applications, human oversight remains essential, particularly where AI outputs may affect customer outcomes or financial provisions.

- **FNOL triage:**
  AI models can read text or voice transcripts at first notification of loss, classify claims and route them to the appropriate team.
- **Fraud detection:**
  ML and DL techniques can identify unusual patterns across claims and policyholder behaviour, signalling potential fraud for human investigation.
- **Operational efficiency:**
  LLMs can assist with drafting letters, summarising files or preparing internal notes, reducing manual workload.
- **Severity and duration prediction:**
  Predictive AI models can estimate ultimate cost and time-to-settlement, supporting reserving and case management.

7.4 Underwriting and Portfolio Management

In underwriting, AI can support the assessment and triage of submissions by extracting relevant information from documents, identifying key risk characteristics and checking alignment with underwriting guidelines. Machine Learning techniques may also be used to analyse portfolio performance, identify emerging trends or highlight segments that deviate from expectations.

AI-enabled tools can help underwriters and actuaries focus their attention on higher-risk or more complex cases, improving efficiency while retaining professional judgement at decision points. Clear controls are required to ensure that such tools are used consistently and that responsibility for decisions remains well defined.

- **Submission analysis:**
  LLMs and DL can extract key information from broker emails, schedules and attachments, aiding initial triage.
- **Appetite checking:**
  RAG can map submissions against underwriting guidelines and highlight areas requiring referral.
- **Portfolio insight:**
  ML models can identify sub-segments of the portfolio that are under-performing or over-performing, guiding corrective actions.

7.5 Capital Modelling and Risk Management

AI may also be applied in capital modelling and risk management, primarily as a decision-support tool. Examples include assisting with scenario analysis, summarising large volumes of risk information, or supporting exploration of risk drivers and sensitivities. However, given the materiality of these areas, robust governance, transparency and Validation are particularly important. Actuaries remain responsible for ensuring that any AI-supported analysis is technically appropriate, well controlled and clearly communicated to stakeholders.

- **Scenario generation:**
  GenAI, guided by risk teams, can help generate plausible stress and scenario narratives, particularly for operational or emerging risks.
- **Policy and risk review:**
  LLMs can assist in reviewing risk policies and related documentation, highlighting inconsistencies or areas requiring further consideration.

7.6 Regulatory Reporting

AI may also support regulatory reporting by helping to structure, summarise and draft narrative content based on actuarial and risk inputs. Language-based AI tools can help prepare sections of reports more efficiently, while maintaining consistency in presentation and messaging. Given the importance of these outputs, careful review and strong governance remain essential. Actuaries remain responsible for ensuring that AI-supported reporting meets regulatory expectations and is accurate, balanced and appropriately evidenced.

- **Reporting:**
  AI can help draft narrative sections of ORSA reports, risk dashboards and Board papers, based on structured inputs from actuarial teams.

Illustrative Example: AI Supporting Reserving Governance and Communication

In many organisations, actuarial reserving relies not only on quantitative models but also on extensive documentation, governance processes and expert judgement. These may include methodology papers, assumption logs, Validation reports, committee minutes and historical commentary. Managing and navigating this material can be time-consuming, particularly where documentation spans multiple years or portfolios.

Language-based AI tools, when appropriately governed, can support reserving teams by acting as decision-support aids rather than decision-makers. For example, such tools may assist in summarising historical reserving rationales, highlighting changes in assumptions over time or helping users locate relevant sections of governance documentation in response to specific queries.

From an actuarial perspective, these tools do not generate reserves or replace expert judgement. Instead, they can improve efficiency, consistency and accessibility of information, helping actuaries focus on analysis, challenge and decision-making while maintaining clear accountability and oversight.

Illustrative example: Claims classification in climate-related events

In practice, insurers often group multiple weather-related impacts under a single cause of loss, such as a "storm event", even though that event may involve different underlying perils, including wind, rainfall, flooding or hail. Distinguishing between these underlying drivers can be challenging due to data complexity, aggregation practices and the volume of information involved.

Research in both academic and industry settings has demonstrated how Machine Learning techniques can support this type of analysis by identifying patterns across large insurance, geographical and meteorological datasets. For example, supervised learning approaches have been used to classify climate-related insurance claims into broad outcome categories (such as no claims, single claims or multiple claims) based on contextual information about location, weather conditions and portfolio characteristics.

From an actuarial perspective, such approaches are not intended to replace professional judgement or established reserving and pricing methodologies. Instead, they can act as decision-support tools, helping actuaries explore complex datasets more efficiently, identify potential drivers of claims experience and inform further investigation and analysis.

7.7 Actuarial Considerations Across AI Applications

The examples in this section are intended to illustrate where AI techniques may be applied in actuarial work, rather than to prescribe how such systems should be designed or implemented. In practice, the appropriateness of using AI depends on the decision context, the materiality of outcomes, data availability and the level of professional judgement required.

When considering AI-enabled approaches, actuaries should reflect on whether the problem genuinely requires additional model complexity or whether traditional actuarial methods remain sufficient and more transparent. In some cases, simpler approaches may be preferable, particularly where decisions are highly regulated, explanations must be robust or data are limited.

AI applications also introduce application-specific risks that require careful consideration. For example, systems used in claims or fraud detection may inadvertently flag legitimate claims at higher rates for certain groups if underlying data reflect historical biases or structural differences. Actuaries should therefore consider how such risks might arise, how they could be detected through monitoring or review and what role human oversight plays in final decision-making.

Across all applications, actuaries should ask clear questions when evaluating AI systems, including what decisions the system supports, how outputs are used in practice, what assumptions and limitations apply and what controls exist to prevent inappropriate reliance on automation. As with other models, appropriate Validation, documentation and governance are essential and the use of AI does not remove the need for professional judgement or accountability.

Detailed application-specific guidance on controls and governance is beyond the scope of this paper and may be developed further through future guidance or practice area publications.

Summary

Across these applications, AI is best viewed as a set of tools that can enhance efficiency, insight and consistency when used appropriately. It does not remove the need for actuarial expertise, professional judgement or accountability. Understanding how AI is applied in practice helps actuaries assess its benefits and limitations and supports informed engagement with AI-enabled processes across the insurance value chain.

8 Practical Applications Across Pensions

As noted in Section 7 the pension examples in this section and the insurance examples in the previous section, are intended to be illustrative rather than exhaustive, highlighting where AI is already commonly used in practice and how it can support actuarial judgement and decision-making. In all cases, the appropriateness of AI use depends on the context, materiality and governance arrangements surrounding the application.

The importance of data and technology were noted by The Pensions Regulator (TPR) in their October 2024 paper[1] titled *"Digital, Data and Technology Strategy – Innovating for better saver outcomes"*. This demonstrates a multi-year focus from TPR to encourage AI and good data practices across the pensions industry.

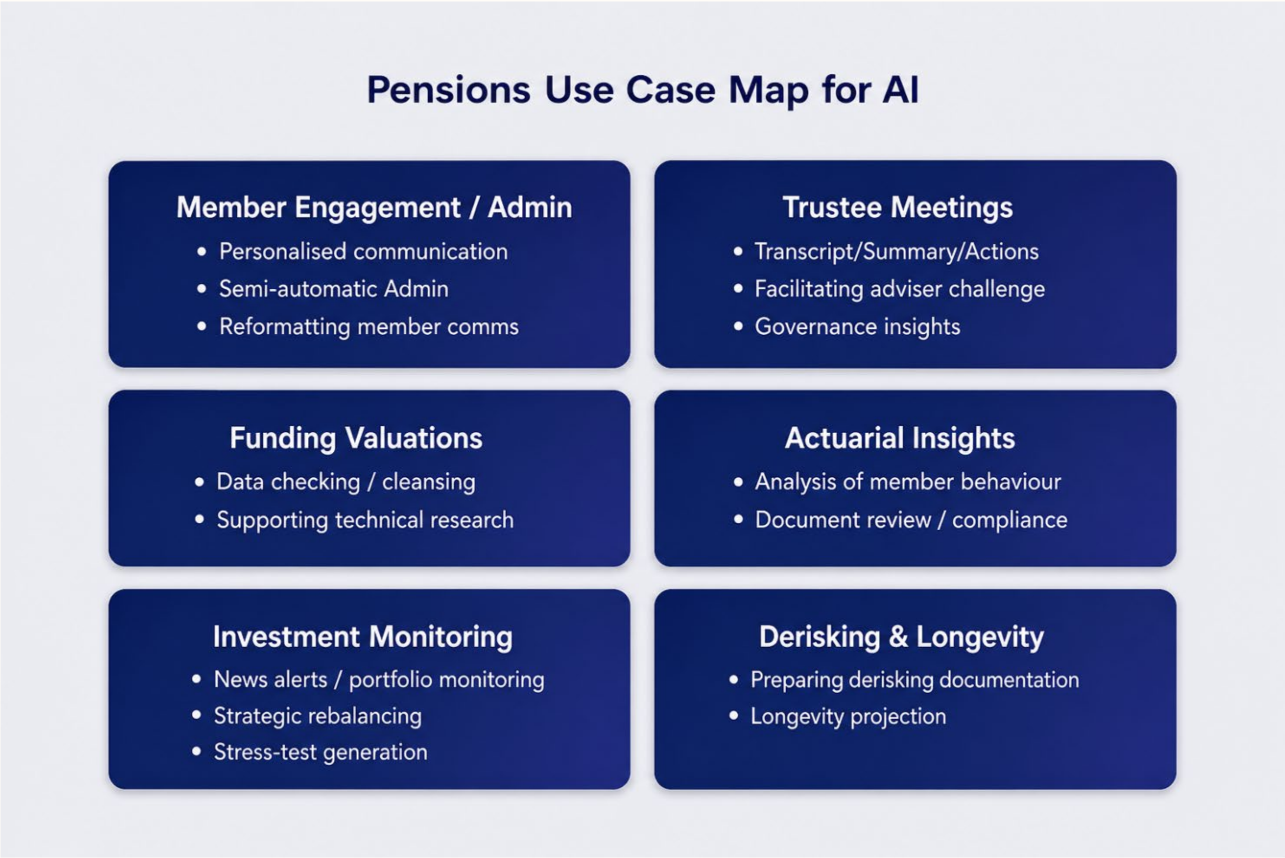

[1] https://www.thepensionsregulator.gov.uk/en/document-library/corporate-information/ddat-strategy

The wide variety of use cases illustrated in figure 4 above shows how AI can already support pension scheme actuaries, administrators, trustees, investment advisers and ultimately pension scheme members / savers.

8.1 Personalised Communication

GenAI can be used to support the creation of personalised communication materials, such as avatar videos, for pension scheme members, for example alongside traditional benefit statements. Video avatars can be used to present information in a consistent and accessible way, with a member's name and selected personal figures and 'nudges' incorporated to improve relevance and engagement.

These approaches also allow messages to be segmented and tailored to specific objectives, such as encouraging increased participation or higher contribution rates. Actuarial involvement can play an important role in this context, helping to ensure that communications are technically accurate, appropriately framed and aligned with relevant professional and ethical standards. This is particularly important where communications may be received by vulnerable individuals, such as elderly members and where careful judgement is required to balance engagement with clarity, fairness and suitability.

Of course, actuaries reviewing communication materials will need to be very conscious of the risks of influencing member behaviour. Communications should be carefully framed to support member engagement in the best interests of the member. In particular, where applicable, communication materials should follow The Pension Regulator's principles for member communications.

8.2 Pensions Administration

In pensions administration, as previously noted, AI is increasingly used to support the handling of incoming member communications, for example by reviewing materials to identify the category of request, such as a change of address or a request for a benefit quotation. These communications may be received through multiple channels, including email, paper correspondence or messaging services within digital applications and AI can help route them efficiently to the appropriate process or team. Agentic AI systems can potentially go further and begin to initiate the administration process (provided appropriate oversight is maintained).

More advanced administration systems may also analyse member sentiment in real time, for example during telephone interactions and generate draft responses for subsequent review by a human operator. Primary responsibility would typically rest with the administrator and actuarial oversight being most relevant where member communications depend on benefit specifications, scheme rules or material calculations. Where such tools are used, it should also be possible to evidence decisions with audit trails (for example, of source materials relied upon and human approvals), so that the use of such tools supports good member outcomes rather than prioritising efficiency alone.

8.3 Refactoring advice into user-friendly formats

In pensions communications, AI may be used to refactor complex or technical written materials into more accessible, member-friendly formats suitable for digital channels. For example, detailed scheme documentation can be transformed into shorter web content, explainer videos or podcast-style summaries that focus on key messages and practical implications for members.

A range of tools exist to support this type of content transformation, including both widely available platforms (e.g. Google's Notebook LM) and proprietary systems developed within administration or advisory environments. Expert actuarial involvement is important to ensure that simplification does not lead to distortion, omission or unintended bias and that the resulting content remains accurate, balanced and aligned with scheme rules and regulatory requirements. Professional oversight also helps ensure that materials are appropriate for the intended audience and support informed member understanding rather than superficial engagement.

8.4 Trustee meeting management

In pensions governance settings, AI may be used to support the documentation of trustee meetings, for example by producing draft notes of discussions and summarising agreed actions. Such tools can help improve efficiency and consistency, particularly for routine or lengthy meetings and may assist trustees and advisers in focusing on key decisions and follow-up items.

AI may also be applied to provide structured feedback on aspects of meeting effectiveness, such as how discussions were facilitated, the extent of participation or how effectively agenda items were progressed.

Actuarial and professional oversight is important in this context to ensure that outputs are accurate, balanced and appropriately interpreted. In particular, care is required to recognise the risk that automated notes or assessments may be incomplete or biased, and that final records and judgements remain the responsibility of the meeting participants.

Actuaries might also be present at meetings when the pension scheme's legal adviser is not present; as a professional, the actuary might be the most appropriate person to query whether appropriate data security checking has been performed (for example, if meetings are being transcribed by AI, where is the associated data being processed and stored?)

8.5 Facilitating *"appropriate challenge"* in Trustee meetings

In pensions governance settings, AI may be used to support non-expert trustees in engaging with and challenging advice presented by technical specialists, such as actuaries. For example, AI tools can help summarise complex advice, highlight key assumptions or risks and suggest areas where further clarification or questioning may be appropriate. More advanced tools may also allow trustees to ask natural-language questions of the underlying advice materials and receive answers that reference the specific source sections relied upon.

Used appropriately, this can help trustees participate more confidently and effectively in decision-making. Actuarial professionalism remains central in this context, both in ensuring that the underlying advice is robust and transparent and in supporting constructive challenge rather than adversarial misunderstanding. Care is required to ensure that AI-generated prompts or summaries do not oversimplify complex issues or create a false sense of assurance, with final judgements resting firmly with trustees and their advisers.

8.6 Data checking and cleaning

In pensions data management, AI may be used to support data manipulation, cleaning and quality assurance activities. For example, AI can assist in developing tools (e.g. a VBA macro in Excel) that programmatically reformat datasets or apply standardised checks, helping to identify anomalies, inconsistencies or missing information within member records, such as unexpected changes in service histories or benefit amounts.

Actuarial involvement plays an important role in this context, helping to ensure that data checks are aligned with scheme rules and benefit structures and that issues are interpreted and prioritised appropriately. While such tools can improve efficiency and consistency, professional judgement remains essential in determining the significance of findings and the appropriate course of remediation.

8.7 Analysis of membership behaviour

In pensions administration and scheme management, AI may be used to analyse employee and membership data to identify patterns or clustering within the membership, including by combining structured records (e.g. contributions, salary, tenure) with unstructured information such as free-text notes, emails or call transcripts. For example, analysis may highlight groups of members at a particular workplace with consistently low contribution rates, or cohorts of employees, such as those who joined during the pandemic, with lower levels of scheme participation. More advanced approaches may also use Predictive AI models to estimate the likelihood of future behaviours (such as opting out, reducing contributions or failing to engage with communications), and may provide a natural-language interface to explore these insights and the drivers behind them.

Such insights can help actuaries and scheme sponsors to better understand behavioural trends and target interventions more effectively. Actuarial involvement is important to ensure that analyses are statistically robust, appropriately interpreted and used in a way that is fair and proportionate. Care is also required to consider data quality, potential bias and the ethical implications of segmenting members, particularly where resulting actions may affect member outcomes.

8.8 Review materials to ensure appropriate for the intended audience

In a pensions context, GenAI may be used to support the review of written communications, helping to ensure that documents are appropriate for their intended audience. For example, member-facing materials can be assessed for clarity and accessibility, with the aim of identifying unnecessary jargon or overly technical language. Similarly, draft client advice documents can be checked for consistency of tone and content, reviewed to ensure a professional style is adopted and cross-checked against relevant checklists.

Such tools may also be used to support compliance and quality assurance processes. GenAI can be applied to review all near-final documents against predefined industry or proprietary checklists, or house-style requirements, highlighting potential omissions or areas of inconsistency. As with other uses of AI in actuarial work, these applications are intended to support, rather than replace, professional judgement and should be accompanied by appropriate human review and governance controls.

8.9 News alerts and portfolio monitoring

In the context of investment management, AI tools can be used in a number of ways. One of the most prevalent is to obtain a daily summary, or real-time news alerts, based on specific market segments or companies. For example, an AI tool can readily scan multiple news sources and provide links directly to preferred sources, e.g. the Financial Times, to stories relevant to, say, the mining sector.

Further, by aggregating such sources over an entire portfolio, the AI tool can monitor specific news and risks in relation to a given investment allocation. This may reveal an aggregation of an emerging risk that had previously not been specifically monitored, e.g. an exposure to a specific currency risk though an accumulation of both direct and indirect exposure.

In particular, AI may be a useful tool to assist with Integrated Risk Management by helping to consider interconnected risks that were not otherwise visibly connected.

8.10 Preparing derisking documentation

As part of any transaction, such as a pension scheme insurance-company buy-in, there is considerable documentation to be prepared. Often this documentation is in a ‘standard form’, covering known subject matter, but it must be tailored to the specific transaction. AI can assist with this process in a number of ways.

Most evidently, documentation can be checked for consistency and compliance by a trained GenAI model. However, some firms have had success going further and, for example, can use an AI tool to prepare a draft benefit specification from a pension scheme's documentation, such as the Trust Deed & Rules.

Summary

As with the use cases in an insurance setting, AI is best viewed as a set of tools that can enhance efficiency, insight and consistency when used appropriately. If you would like further information about any of the use cases mentioned in this section, please contact the authors.

It is important to note that AI does not remove the need for actuarial expertise, professional judgement or accountability. Understanding how AI is applied in practice helps actuaries assess its benefits and limitations and supports informed engagement with AI-enabled processes in the pensions and saving environment.

9 Risks, Limitations and Challenges in AI Adoption

While AI offers significant opportunities to support actuarial work, it also introduces a range of risks and limitations that must be understood and managed. In this section we highlight some of the key risks, although this should not be regarded as an exhaustive list.

Many of these risks are not entirely new, but they may be amplified by the scale, complexity and perceived authority of AI-enabled systems. Recognising these challenges is essential for responsible adoption.

9.1 Data Quality and Representativeness

AI systems are highly dependent on the quality, relevance and representativeness of the data on which they are trained and operated. Incomplete, biased or poorly understood data can lead to misleading outputs, regardless of the sophistication of the underlying techniques. Changes in data definitions, collection processes or portfolio composition can further undermine model reliability. For actuaries, this reinforces the importance of understanding data provenance, limitations and suitability for purpose. Data issues that might be manageable in simpler models can have more pronounced effects in complex AI systems.

Many AI-related risks can be viewed through familiar actuarial concepts such as parameter uncertainty, model uncertainty, epistemic uncertainty and changing external conditions over time. While these are not new challenges, the scale and opacity of some AI systems can make them harder to identify and monitor. This reinforces the importance of established actuarial disciplines such as sensitivity testing, Validation, monitoring and professional judgement.

9.2 Bias and Fairness

AI models trained on historical data may reflect or amplify existing biases present in that data. In actuarial contexts, this raises significant regulatory, ethical and reputational concerns, particularly where AI outputs influence individual outcomes of customers, employees, or pension scheme members.

A key challenge is that bias may arise even when protected characteristics (such as age, gender, ethnicity or disability) are not explicitly included as model inputs. Proxy discrimination can occur where other variables correlate strongly with protected characteristics, leading to systematic differences in outcomes across groups. This is particularly relevant in complex Machine Learning models, where relationships between variables may be less transparent.

There is also an inherent tension between actuarial risk segmentation and fairness considerations. For example, while risk differentiation is fundamental to insurance pricing and underwriting, regulatory and ethical expectations require that outcomes remain reasonable, proportionate and justifiable. Actuaries play a critical role in assessing where this balance lies and in ensuring that AI-supported decisions remain defensible.

From a practical perspective, addressing fairness requires more than a one-off review. It may involve testing outcomes across different groups, monitoring for changes over time and considering whether adjustments or constraints are needed to mitigate unintended impacts. Clear documentation of assumptions, limitations and trade-offs is essential, particularly when AI systems are used as part of regulated decision-making processes.

9.3 Hallucinations and Incorrect Outputs

Language-based AI systems, including Large Language Models, can generate outputs that appear fluent, confident and plausible but are factually incorrect, incomplete or unsupported by evidence. This phenomenon is commonly referred to as *"Hallucination".* Unlike traditional calculation errors, Hallucinations may be difficult to detect because the output is presented in natural language and often aligns with the user's expectations.

In actuarial contexts, Hallucinations are most likely to arise when LLMs are used for tasks such as summarising documentation, drafting reports, answering technical queries or interpreting regulatory or methodological material. Without appropriate controls, such systems may fabricate details, misstate assumptions or invent references that do not exist. The risk is therefore not limited to technical modelling errors, but extends to communication, governance and decision-support processes.

Techniques such as Retrieval-Augmented Generation can help reduce these risks by grounding outputs in approved source material. However, they do not eliminate the possibility of incorrect or misleading responses, particularly where source documents are incomplete, outdated or ambiguously phrased. As a result, AI-generated text should be treated as a starting point rather than an authoritative source. Actuaries should apply professional scepticism and ensure appropriate human review, especially where outputs inform regulatory submissions, Board materials or customer-facing communications.

9.4 Bad Actors and the Alignment Problem

AI technologies can be exploited by bad actors to support malicious or unlawful activities, including cybercrime, fraud, misinformation and social engineering. The same capabilities that make AI systems powerful for legitimate purposes, such as generating realistic text, automating interactions or analysing large datasets, can also be repurposed to scale harmful activity at relatively low cost.

Examples include the use of Generative AI (GenAI) to produce convincing phishing communications, automate reconnaissance of systems or assist in the creation of malicious code.

These risks are not confined to highly sophisticated adversaries. The increasing accessibility of AI tools lowers the barrier to entry for misuse, enabling a broader range of actors to undertake activities that previously required significant technical expertise. This creates challenges for organisations seeking to protect data, systems and customers, particularly where AI-generated outputs are difficult to distinguish from legitimate human activity.

Closely related to the risk of bad actors is the alignment problem. Alignment refers to the challenge of ensuring that AI systems behave in ways that are consistent with intended objectives, ethical norms and regulatory expectations. In practice, there is often no clear boundary between acceptable and unacceptable use. Systems designed to be flexible and helpful may be deliberately or inadvertently prompted in ways that produce harmful, misleading or inappropriate outputs. Determining where legitimate use ends and misuse begins is therefore non-trivial, particularly in open-ended or user-driven systems.

One tangible example of this dilemma was whether having an LLM capable of amending the clothing of people in a picture was appropriate. Some commentators thought it was appropriate to have an image adjusted to change clothing from, for example, *"sweatshirts to dress suits"* – in the current environment this was generally deemed appropriate. However, *"sweatshirts to bikinis"* wasn't deemed appropriate. The more libertarian-positioned AI houses supported the more extreme views on the grounds of *"free speech"* and this debate and similar debates, will no doubt continue for some time. There have already been many debates on this subject, and the discussion will no doubt continue into the future.

Misalignment can arise even in the absence of malicious intent. Poorly specified objectives, inadequate constraints or insufficient testing can result in behaviour that is technically consistent with a system's design but misaligned with organisational values or professional standards. In more advanced or autonomous systems, this risk may be amplified by complex interactions between components or by emergent behaviour that was not anticipated at the design stage.

Both bad actor misuse and misalignment represent significant risks for AI development and deployment. Mitigating these risks requires a combination of technical controls, such as access restrictions, monitoring and guardrails and organisational measures, including clear acceptable-use policies, escalation processes and accountability structures. For actuaries, this reinforces the importance of viewing AI not solely as a technical tool but as part of a broader socio-technical system, where governance, incentives and human judgement play a critical role in ensuring responsible use.

9.5 Model Performance Degradation and Data Drift

AI model performance can deteriorate over time as underlying patterns in the data change. This may be due to shifts in customer behaviour, economic conditions, legal environments or operational processes. This phenomenon, often described as performance degradation due to Data Drift, can be difficult to detect without appropriate monitoring.

Regular performance assessment, clear thresholds for intervention and defined retraining or recalibration processes are therefore essential components of AI governance.

9.6 Over-reliance on Automation

The apparent sophistication of AI systems can create a risk of over-reliance, where users place undue confidence in model outputs or reduce the level of professional challenge applied to them. This risk is particularly acute when AI systems are presented through polished interfaces or produce outputs in fluent natural language. AI systems should be designed and used as decision-support tools rather than decision-makers. Clear guidance on appropriate use, combined with training and cultural reinforcement of professional scepticism, is critical.

9.7 Interpretability and Explainability

The terms *interpretability* and *Explainability* are often used interchangeably but can refer to different concepts, particularly in regulatory and governance contexts. Interpretability generally relates to understanding how a model operates internally, for example, how input variables influence outputs. Explainability, by contrast, focuses on providing understandable factors for specific model outputs or decisions to stakeholders such as regulators, Boards, customers or internal users.

Different AI technologies present different Explainability challenges. For traditional statistical or Machine Learning models, interpretability may be supported through variable importance measures, sensitivity analysis or simplified representations. For language-based models, Explainability is typically more limited and may instead rely on transparency around data sources, prompt design, constraints and usage controls. In more complex or agent-based systems, Explainability may involve understanding task sequencing, decision logic and escalation paths rather than causal relationships.

In regulated environments, the appropriate level of interpretability and Explainability should be proportionate to the materiality of the decisions supported by the AI system. Actuaries should therefore be cautious when claims of *"Explainable AI"* (XAI) are made and seek clarity on what is meant in practice. This might include, for example, investigating what limitations apply, what underlying assumptions have been made and whether the explanations provided are sufficient for the intended use. Where full transparency is not feasible, compensating controls, such as restricted use, enhanced monitoring or human review, may be necessary.

9.8 Governance Complexity

AI systems combine data, models, software and human processes. This creates new forms of model risk and operational risk. Governance frameworks need to be updated to cover AI specific issues such as prompt management, RAG source control, agent behaviour and emergent properties.

Summary

These risks do not imply that AI should be avoided, but they do highlight the need for careful design, governance and oversight. Many of the skills required to manage these challenges, including professional scepticism, Validation, documentation and clear communication are core actuarial competencies. Applying them consistently to AI-enabled systems helps ensure that AI supports actuarial work in a reliable and responsible manner. In particular, issues of fairness, Explainability and ongoing behavioural monitoring require careful consideration, as AI systems may introduce new risks even where traditional actuarial controls are already in place.

10 Responsible AI & Governance for Actuaries

The use of AI in actuarial work does not remove the need for governance; rather, it increases its importance. AI systems introduce new sources of complexity, opacity and operational risk, but they also draw on many of the same principles that already underpin actuarial model governance. As a result, actuaries are well placed to contribute to the responsible oversight of AI-enabled systems.

Responsible AI governance begins with clear purpose and accountability. The intended use of an AI system should be explicitly defined, along with the decisions it is designed to support. Roles and responsibilities must be clearly established, including who is accountable for approving, monitoring and challenging AI outputs. This is particularly important where AI systems influence material financial outcomes or customers'/savers' decisions.

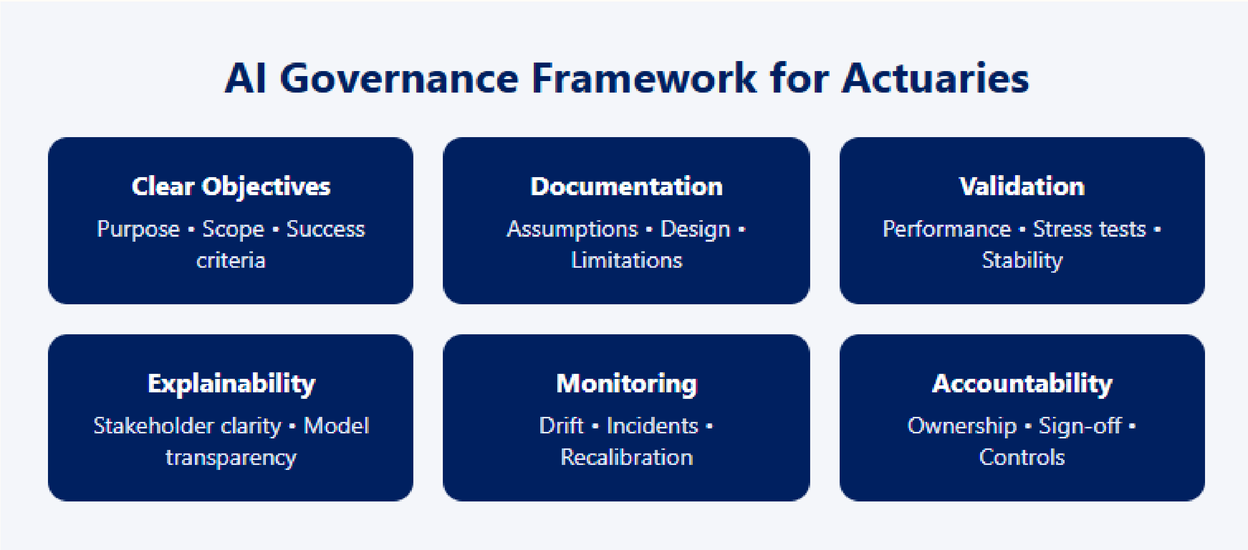

The wide variety of use cases illustrated in figure 5 above shows AI frameworks from a Governance perspective. In this section we will define each of the above and then dive into the detail of AI Governance for actuaries.

Key generic governance principles, based on the actuarial control cycle, include:

- **Clear objectives:**
  The purpose of an AI system should be clearly stated, including the decisions or processes it is intended to support and the boundaries within which it should operate. Success measures, constraints and intended users should be defined at the outset. This helps avoid inappropriate use, misunderstanding of outputs or misalignment with business and regulatory expectations.

- **Documentation:**
  AI systems should be documented in a manner proportionate to their materiality and complexity. Documentation should cover data sources, key assumptions, limitations, intended use, Validation results and known failure modes. Where a system includes language-based components, documentation should also record any relevant design choices such as prompt structures, retrieval sources and controls over knowledge bases. Where appropriate the documentation requirements of TAS-M (Technical Actuarial Standard M: Modelling) should be considered.
- **Validation and testing:**
  AI systems require robust Validation prior to deployment and periodic reValidation thereafter. Testing should include out-of-sample performance assessment, sensitivity testing and checks for stability across relevant segments of the portfolio. Where appropriate, Validation should also consider bias and fairness and should include challenge against benchmarks or alternative methods. Validation should be supported by clear sign-off processes, reflecting the system's use and materiality.
- **Explainability and interpretability:**
  The level of understanding and explanation required should be proportionate to the impact of the decisions being supported. In regulated environments, stakeholders may require a clear explanation of model behaviour, key drivers, outputs and limitations, even where full technical transparency is not feasible. Where this is limited, compensating controls may be required, such as constrained use, additional human review or simplified supporting explanations.
- **Monitoring:**
  Once deployed, AI systems should be monitored to ensure continued suitability and performance. Monitoring should consider performance degradation due to Data Drift, changes in portfolio mix, operational changes and shifts in external conditions. Organisations should define thresholds that trigger investigation, recalibration, retraining or withdrawal of a model.
- **Accountability:**
  Clear accountability must be established for AI-enabled decisions. AI tools may support users, but responsibility for decisions remains with people and the governance structures of the organisation. Roles should be defined for model ownership, approval, oversight, incident management and escalation. This is particularly important where AI outputs influence customer outcomes or material financial decisions.

Over-arching all of the above are the actuarial profession's commitment to strong ethics. In particular, AI adoption should always reflect the actuarial profession's ethical standards and public interest obligations. This includes considering fairness, proportionality, transparency and the potential consequences of automated or AI-supported decisions. Actuaries should apply professional scepticism and ensure that AI is used in a way that supports sound judgement, clear communication and responsible treatment of customers and stakeholders.

These principles should apply whether AI is developed internally or procured from vendors.

10.1 Governing AI Systems as Interconnected Components

Unlike traditional actuarial models, many AI-enabled systems are composed of multiple interdependent components, such as data pipelines, feature stores, retrieval mechanisms, application interfaces and downstream decision logic. The behaviour of the overall system depends not only on individual components, but also on how they interact.

From a governance perspective, actuaries should therefore consider the AI system as a whole rather than focusing solely on individual component models. This includes understanding where data enters the system, how information flows between components, which elements are automated and where human judgement or intervention is expected to be needed. Clear ownership and accountability should be defined not just for component models, but for the full end-to-end process.

10.2 Explainability in an AI Context

Explainability in AI does not have a single meaning and varies depending on the type of system involved. For traditional Machine Learning models, Explainability may focus on understanding how input variables influence outputs. For language-based models, Explainability is more limited and often relates to transparency around training data sources, prompt design and output constraints rather than causal interpretation. In Agentic systems, Explainability may involve understanding how tasks are sequenced and how decisions are delegated across components.

Actuaries should therefore be cautious when claims of *"Explainable AI"* are made and seek clarity on what is meant in practice. The key question is whether the level of explanation provided is sufficient for the intended use, the materiality of decisions supported and the regulatory and professional standards that apply.

10.3 Monitoring and Ongoing Oversight of AI Behaviour

Monitoring AI-enabled systems extends beyond tracking traditional performance metrics. For systems that incorporate language models or retrieval mechanisms, additional risks may arise, including changes in behaviour due to data updates, prompt modifications or interactions between components. In more autonomous or Agentic systems, unexpected or emergent behaviours may occur when systems operate over time or across multiple steps.

Effective governance therefore requires ongoing oversight that considers not only statistical performance, but also changes in system behaviour, reliance on automation and the conditions under which human review or intervention is required. Monitoring arrangements should be proportionate to the complexity and materiality of the system and should be revisited as systems evolve.

10.4 Evaluating Vendor Claims and Third-Party Systems

Where AI systems are sourced from vendors or embedded within third-party platforms, actuaries should not rely solely on high-level assurances regarding fairness, bias or Explainability. Instead, governance should include clear questions about how such claims are substantiated, what evidence is available and what limitations apply.

This may involve understanding what aspects of the system are configurable, what documentation is provided, how updates are managed and how responsibilities are shared between the organisation and the vendor. As with other models, the use of third-party AI does not remove the need for actuarial oversight or professional accountability.

Summary

Responsible AI governance builds on familiar actuarial principles: clear objectives, strong documentation, appropriate Validation, proportionate Explainability, ongoing monitoring, defined accountability and ethical professionalism. By applying these consistently, actuaries can help ensure AI-enabled systems are robust, trustworthy and aligned with both regulatory expectations and professional standards. Effective AI governance therefore requires not only established actuarial model controls, but also consideration of system interactions, Explainability expectations, ongoing behavioural monitoring and scrutiny of third-party claims.

11 Emerging Technologies Actuaries Should Be Aware Of

The AI landscape continues to evolve rapidly and new techniques, tools and infrastructure continue to emerge. While this paper focuses primarily on established foundations and current actuarial applications, it is helpful for actuaries to maintain a high-level awareness of developments that may become more relevant over time. The technologies described in this section are included to support horizon scanning and contextual understanding, rather than to suggest immediate adoption.

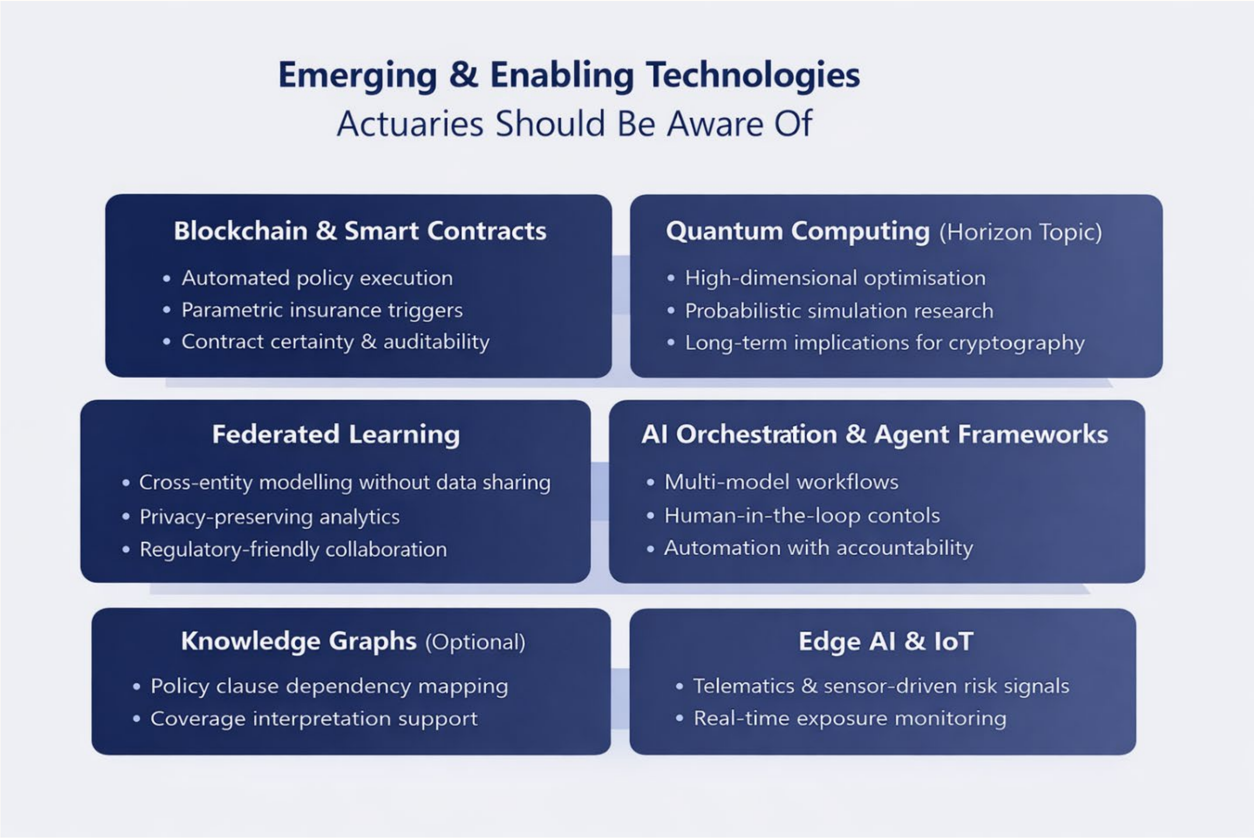

Figure 6 above provides a high-level overview of selected emerging and enabling technologies that actuaries may encounter or should be aware of in future actuarial and insurance applications.

11.1 Blockchain & Smart Contracts

Blockchain and Smart contracts are enabling technologies that can support trusted record-keeping and automated execution across multiple parties. In insurance, the most frequently cited application is parametric cover, where a Smart contract may initiate a payment when an agreed external trigger is met, subject to appropriate governance over the trigger data source. These tools are not required for most actuarial modelling, but actuaries may encounter them where products rely on event triggers, automated settlement logic or multi-party data reconciliation.

Examples:
• Parametric insurance contracts, such as catastrophe or weather-based covers, may use Smart contracts to automate pay-outs when predefined external triggers are met, for example wind speed, rainfall or earthquake magnitude. Once the trigger is verified by an agreed data source, the contract logic may initiate settlement without the need for a traditional claims assessment.

• In more complex insurance or reinsurance arrangements involving multiple parties, blockchain-based ledgers may be used to record agreed claim events, timestamps and settlement milestones. Smart contracts may then enforce predefined payment logic once conditions are met, reducing reconciliation effort across insurers, reinsurers and intermediaries.

• Smart contracts may also be used to track policy endorsements or coverage changes over time, helping ensure that pricing, exposure and claims handling reflect the correct contract state at each point in time.

11.2 Quantum Computing

Quantum computing is an emerging computational approach that uses quantum-mechanical principles to process certain classes of problems in ways that may differ materially from classical computing. While current practical actuarial applications remain limited, potential future relevance may include optimisation, scenario modelling and complex simulation. Quantum computing is best regarded at present as a horizon topic rather than a near-term actuarial tool.

Actuaries do not need detailed technical expertise at this stage, but awareness may still be useful given the possible longer-term implications for modelling, analytics and cryptographic infrastructure.

Examples:
• Research into high-dimensional optimisation problems that may be difficult to solve efficiently using classical computing approaches.
• Exploration of probabilistic simulation techniques that may, over time, influence complex risk modelling.
• Serious implications for cryptography and secure digital infrastructure relevant to financial services. One scenario is that Quantum Computers will be able to decrypt messages that are current end-to-end encrypted by, say, 2030. It is believed that bad actors are already *"harvesting"* encrypted data today, with the aim of decrypting the data at some point in the future.

11.3 Federated Learning

Federated learning refers to approaches in which models are trained across multiple datasets without centralising the underlying data. Instead, learning occurs locally, with model updates aggregated centrally. This can help address data privacy, confidentiality and regulatory constraints.

Such techniques may become relevant where organisations wish to collaborate on modelling while retaining control over sensitive data. For actuaries, this raises considerations around Validation, consistency, data representativeness and governance across distributed environments.

Use case example: Google Gboard

A well-known real-world example is Google's Gboard keyboard, where federated learning has been used to improve language models across many users' devices without centralising the underlying typed text. This illustrates the broader principle that a model can be improved using experience held across multiple locations while keeping the underlying data local. In future actuarial settings, a similar concept could become relevant where organisations wish to collaborate on modelling without sharing raw policyholder or member data. Actuaries can add value by ensuring appropriate safeguards are applied to the federated learning process, as well as benefiting from data accumulated via such data pooling.

11.4 AI Orchestration and Agent Frameworks

As AI systems become more complex, tools are emerging to coordinate multiple models, data sources and processes into coherent workflows. These orchestration and agent frameworks manage sequencing, routing, decision logic and integration with existing systems. In practice, they may combine language models with more targeted techniques such as document classification, information extraction, workflow rules, retrieval components and human review checkpoints.

From an actuarial perspective, these frameworks may affect how models are combined, how outputs move through business processes and where governance controls need to be applied. They are also relevant in settings where different AI components work together across structured and unstructured inputs, including text, tables, forms, images and other operational data.

Examples:

- End-to-end underwriting or claims workflows that chain together several AI components, together with human checkpoints and approval controls.
- Pensions administration systems that process incoming mail, redirect it to appropriate sub-processes, create a draft response, propose a database change and then present the full package to a human operator for Validation.
- Document-handling workflows that combine targeted Natural Language Processing (NLP)

techniques, retrieval tools and language models to classify documents, extract key information and generate summaries or draft outputs.

Actuaries do not need to be experts in each of these technologies, but awareness is valuable in understanding the direction of travel and the opportunities and risks that may arise as AI becomes more embedded in operational processes.

11.5 Knowledge Graphs

Knowledge graphs represent information as entities and relationships, allowing complex connections to be stored and queried explicitly. In insurance settings, they may be used to model relationships between policies, coverages, exclusions, jurisdictions, claims or counterparties.

When combined with language-based AI systems, knowledge graphs may support more structured reasoning, reduce ambiguity and improve consistency in how relationships are interpreted. For actuaries, they may become relevant in areas such as policy interpretation, exposure mapping, risk aggregation or governance support. While not entirely new as a concept, they remain a useful enabling technology to be aware of in more advanced AI-enabled environments.

Examples:
• Mapping complex coverage dependencies across programmes and layers.
• Supporting users in understanding how changes in one clause may affect others.
• Linking related policy, claims and exposure information to support more structured querying and interpretation.

11.6 Edge AI and IoT

Edge AI involves running models on devices or sensors close to where data are generated, rather than processing everything centrally. In insurance, this may include telematics devices, sensors in property risks, industrial monitoring equipment or wearable technologies. These systems may generate continuous or near-real-time signals that influence operational monitoring and risk assessment.

These developments may affect the nature, volume and timeliness of data available for pricing, reserving and risk monitoring. Actuaries should be aware of how such data sources could influence portfolio dynamics, segmentation and assumptions, even where they are not directly involved in system design.

Examples:
• Real-time exposure or risk monitoring in motor or property portfolios.

- Usage-based insurance products informed by telematics or behavioural signals.
- Sensor-driven risk alerts that support operational intervention or portfolio monitoring.

Summary

Emerging and enabling technologies continue to expand the range of possible applications relevant to actuarial work. Not all developments will be immediately relevant and not all will prove appropriate in actuarial contexts. However, awareness of these trends can help actuaries anticipate future change, ask informed questions and assess potential implications for data, models, governance and professional judgement. As with more established AI techniques, any adoption should be guided by proportionality, clear accountability and robust oversight.

12 What Actuaries Should Prioritise

As AI becomes more embedded across insurance and financial services, actuaries do not need to become technologists or software engineers. However, they do need to develop a set of capabilities that enable them to engage confidently with AI-enabled systems, assess their appropriateness and exercise professional judgement. The priorities outlined below reflect areas where actuarial skills are most valuable in an AI-supported environment.

12.1 Data Literacy

A strong understanding of data remains fundamental. Actuaries should be able to assess data quality, relevance and limitations, including how data are collected, processed and transformed before being used in models. This includes awareness of data provenance, potential biases and changes in data definitions over time. Data literacy enables actuaries to identify risks early and to challenge AI outputs appropriately:

- Interpreting data dictionaries and schemas.
- Identifying common data issues and their impact on models.
- Working collaboratively with data engineers and data scientists.

12.2 Model Understanding and Explainability

Actuaries should aim to understand how AI models behave at a conceptual level, even where full technical detail is not required. This includes knowing what inputs drive outputs, where models perform well and where they may be unreliable. Understanding limitations, such as sensitivity to Data Drift, bias or sparse data is essential for informed use and oversight.

Effective communication is also critical when AI is used to support decision-making. Actuaries should be able to explain, in clear and proportionate terms, how AI-supported analyses have informed outcomes, including the associated assumptions, uncertainties and limitations. This applies equally to communication with Boards, regulators, colleagues and other stakeholders.

- What inputs the model uses.
- What outputs it produces and how they are used.
- How performance is measured.
- Where the model is likely to perform well or poorly.

This is essential for informed challenge and communication with Boards, regulators and other stakeholders.

12.3 Governance and Model Risk

AI systems introduce new forms of model and operational risk, but they also sit naturally within existing actuarial governance frameworks. Actuaries should continue to play a leading role in defining Validation standards, documentation requirements, monitoring processes and escalation procedures. Applying established model risk management principles consistently to AI-enabled systems helps ensure robustness and accountability:

- Cover both traditional and AI models coherently.
- Include clear criteria for model approval and retirement.
- Define how models are monitored and revalidated.

12.4 Professional Scepticism

Core actuarial skills, scepticism, prudence and careful examination of assumptions, are more important than ever. AI outputs should be challenged and contextualised, particularly where they conflict with experience, intuition or other evidence. Actuaries should guard against over-reliance on automation and remain alert to situations where human judgement is required:

- "What is this model actually doing?"
- "What assumptions is it making implicitly?"
- "How might this go wrong?"

12.5 Communication and Collaboration

Explaining AI-enabled processes in clear, non-technical language is a critical skill. Actuaries are often the bridge between technical teams (data science, engineering) and decision-makers (Boards, regulators, clients). Being able to communicate benefits, risks and limitations is essential.

In particular, AI initiatives are typically multidisciplinary, involving data scientists, engineers, technology teams, legal and compliance functions and business users. Actuaries should be prepared to collaborate effectively across these disciplines, contributing domain knowledge, risk insight and governance expertise. Clear communication across disciplines helps ensure that AI systems are designed and used appropriately:

- Data scientists and ML engineers.
- Software developers and IT teams.
- Legal, compliance and risk experts.
- Operational teams in underwriting, claims and finance.

Collaboration does not dilute actuarial identity; it amplifies it in new contexts. Actuaries should be acutely aware of the danger of professional “overreach”, facilitated by AI tools; we should respect the value brought by other professionals in the same way as we expect to be respected.

12.6 Lifelong Learning

The AI landscape continues to evolve. Actuaries should engage in ongoing learning to maintain awareness of relevant developments, regulatory expectations and emerging risks. This does not require deep technical specialisation, but it does require a willingness to update knowledge and adapt professional practice over time. The IFoA is available to assist with this, regularly providing events and resources to help actuaries remain up to date.

Summary

By focusing on these priorities, actuaries can remain effective and influential in AI-enabled environments. The profession’s value lies not in replicating technical capabilities, but in providing judgement, oversight and clear communication. These strengths are essential to ensuring that AI is used responsibly, proportionately and in the public interest.

13 The Future: AI-Accelerated Actuarial Practice

AI is reshaping what actuaries do and how they do it. This is illustrated in figure 7 below:

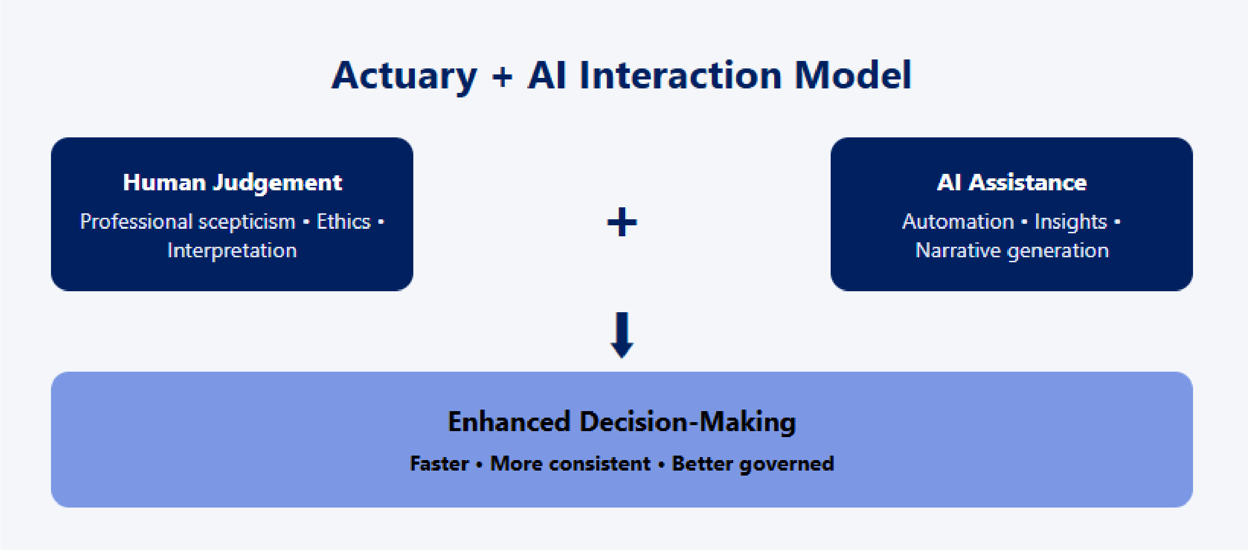

In an AI-accelerated future:

- Routine data preparation and documentation will be increasingly automated.
- Models will be easier to run, compare and embed within business workflows.
- Actuaries will spend more time on design, oversight, communication and strategic decision-making.
- Budgeting and scheduling time for developers to create a "wireframe" demonstration of a software tool can now be replaced by an AI-built working "mock-up" application
- Interaction with models and data will become more conversational and interactive (e.g. "asking questions" of a model or a portfolio).

The core actuarial value proposition, which is robust quantification of risk, professional judgement and clear communication, remains unchanged. AI, when adopted responsibly, is a powerful set of tools to support that work.

14 Conclusion — Clarity, Capability and Confidence

Artificial Intelligence is increasingly influencing how actuarial work is performed, communicated and governed. While the underlying technologies may evolve rapidly, the professional responsibilities associated with their use remain familiar. Actuaries are not required to become technologists, but they do need sufficient understanding to engage confidently with AI-enabled systems, assess their suitability and apply appropriate professional judgement.

This paper has sought to provide a clear, non-technical foundation for understanding the AI landscape, focusing on concepts and applications most relevant to actuarial practice. By distinguishing between different AI technologies, outlining common areas of use and highlighting associated risks and governance considerations, the paper aims to support informed professional discussion rather than prescribe specific tools or solutions.

Many of the challenges associated with AI, including data quality, model risk, Explainability and ethical considerations, align closely with established actuarial principles. Applying these principles consistently to AI-enabled systems helps ensure that such technologies are used responsibly, proportionately and in the public interest.

AI adoption within insurance, pensions and other actuarial disciplines will continue to develop and no single document can be comprehensive or definitive. This paper should therefore be viewed as a starting point rather than an endpoint. Ongoing professional engagement, learning and dialogue will be essential as the actuarial profession navigates the opportunities and risks presented by AI in the years ahead.

By building a shared, non-technical understanding of AI foundations, actuaries can engage constructively with AI initiatives, provide informed challenge and oversight, support Boards and regulators in understanding implications, identify where AI can enhance existing actuarial work and contribute to the design of governance frameworks that keep human judgement central.

The AIET Practice Board will continue to support members through guidance, education and dialogue. This paper is intended as a baseline reference point, a starting framework from which actuaries can deepen their understanding, ask the right questions and help shape the future of AI-enabled actuarial practice with clarity, capability and confidence.

Scope Clarification

This paper focuses on the AI technologies most relevant to current actuarial practice. For clarity and accessibility, several emerging tools and methods are not explored in detail. It is intended that the Emerging Technology Pillar of the AIET practice board will continue to publish horizon scans and insight papers to keep the profession informed as the technology landscape evolves.

15 Glossary of Key Terms

Where available, definitions are aligned to the Alan Turing Institute's (ATI) "Data science and AI glossary" which can be accessed in full here[2]. Terms not covered there (including actuarial governance terms and newer GenAI concepts) are defined for the purposes of this paper.

Terms aligned to the Alan Turing Institute (ATI) glossary

Artificial Intelligence (AI)
The design and study of machines that can perform tasks that previously required human (or other biological) intelligence (e.g. reasoning, learning, decision-making and communication).

Algorithmic Bias
Algorithmic bias in this paper refers to unfair or unjustified differences in model outcomes across individuals or groups, rather than the technical statistical meaning of bias used in Machine Learning theory.

Machine Learning (ML)
A field of AI in which algorithms learn patterns from sample data and apply what they learn to new data to produce useful outputs (e.g. predictions or translations).

Deep Learning (DL)
A form of Machine Learning using neural networks (often with multiple layers) that automatically recognise patterns and produce outputs such as predictions or evidence for decisions.

[2] https://www.turing.ac.uk/news/data-science-and-ai-glossary

Generative AI (GenAI)
AI systems that generate content (e.g. text, images, audio or video) in response to prompts, using Machine Learning to create outputs resembling the data they were trained on.

Large Language Model (LLM)
A type of foundation model trained on very large amounts of text to perform language tasks and generate text; widely used in modern chatbots.

Natural Language Processing (NLP)
A field of AI focused on analysing or generating human language (speech and text), supporting applications such as translation, speech recognition and sentiment analysis.

Synthetic Data
Artificially generated data that can retain statistical properties of real data while removing identifying information; often used where privacy is important.

Foundation Model
A Machine Learning model trained on a vast amount of data that can be adapted for many applications; LLMs are a common example.

Additional terms used in this paper (not defined in the ATI glossary)

Agentic AI / Agentic systems
AI systems designed to plan and carry out sequences of actions toward a defined objective, often combining models, retrieval and tools, operating within constraints and requiring appropriate oversight. May include access to external systems, such as via a web browser or a synthetic PC.

Application Programming Interface (API)
A standardised way for one piece of software to request services from another. In an AI context, calling a model *"via an API"* means sending prompts to a provider (e.g. OpenAI, Anthropic, Google) over the internet and receiving responses, rather than running the model yourself.

Data Drift
A change over time in the statistical properties of incoming data compared with what a model was trained on, which can contribute to degraded performance and should be monitored.

Explainability / Explainable AI / XAI
The extent to which a model's outputs, behaviour and limitations can be understood and communicated to appropriate stakeholders, proportionate to the materiality of use. Actuaries should be aware that there is some ambiguity in the use of this term: some commentators use this term to mean *"we can explain how the computational process got to this answer"*, others use it to

mean *"we can approximate the key underlying factors that give rise to this output"*. The term can sometimes be misunderstood to me *"interpretable"* (see Section 9.7 for more details). Sadly, many vendor claims of Explainability are only approximations or post-hoc interpretations, rather than true causal explanations.

Hallucination
Where the underlying algorithm of a language model generates plausible-sounding content that is incorrect or unsupported and hence appears *"fabricated"*.

Model Governance
The processes, controls and accountability structures ensuring models are appropriate for their intended use, Validated, documented, monitored and periodically reviewed.

Retrieval-Augmented Generation (RAG)
A pattern that retrieves relevant information from approved sources (documents/data) and uses it to ground generated outputs, reducing unsupported responses. The key advantage of RAG for an actuary is that it lets an LLM specifically answer from *your* documents (regulatory texts, internal methodology notes, client data) rather than from its training data alone.

SaaS
Software as a Service (renting software capability, in contrast to selling a piece of software for a fixed one-off sum).

Validation
Assessment of whether a model is adequate for its intended purpose, including testing, sensitivity/robustness checks and review of assumptions and limitations.

Predictive AI
Predictive AI refers to systems that use historical data to forecast outcomes or classify new cases. For example, estimating lapse rates or flagging fraudulent claims. This encompasses most traditional Machine Learning and contrasts with Generative AI and Agentic AI. The bulk of AI work that has involved actuaries to date has been Predictive AI, as it aligns naturally with existing modelling discipline and Validation practice. (It is also worth noting that some commentators regard the term "AI" to exclude Predictive AI.)

Quantum Computing
A computing approach based on quantum mechanics; not an AI technology in itself and currently not a mainstream tool in actuarial practice.

Self-Hosted / Open-Weight Models
Large Language Models whose weights are publicly released (e.g. Llama, Mistral, DeepSeek) and which can be run on a firm's own infrastructure rather than via a third-party API. For work involving confidential client data, member records or regulated information, self-hosting may be the solution to maintaining confidentiality.

Smart contract
A self-executing digital agreement in which predefined rules are automatically carried out when specified conditions are met, typically on a blockchain.

16 References:

- **Institute and Faculty of Actuaries (IFoA)** (2023).
  Ethical Standards and Professional Principles.
  Available at: https://actuaries.org.uk

- **Institute and Faculty of Actuaries (IFoA)** (2024).
  Model Risk Management and Governance: Principles for Actuaries.
  Available at: https://actuaries.org.uk

- **Bank of England and Prudential Regulation Authority (PRA)** (2022).
  Model Risk Management Principles for Financial Institutions.
  Available at: https://www.bankofengland.co.uk

- **Organisation for Economic Co-operation and Development (OECD)** (2019).
  OECD Principles on Artificial Intelligence.
  Available at: https://www.oecd.org/ai

- **National Institute of Standards and Technology (NIST)** (2023).
  Artificial Intelligence Risk Management Framework (AI RMF 1.0).
  Available at: https://www.nist.gov/itl/ai-risk-management-framework

- **European Union** (2024).
  Artificial Intelligence Act (AI Act).
  Available at: https://artificialintelligenceact.eu

- **UK Government** (2023).
  A Pro-Innovation Approach to AI Regulation.
  Department for Science, Innovation and Technology.
  Available at: https://www.gov.uk/government/publications

- **The Pensions Regulator (TPR)** (2024).
  Digital, Data and Technology Strategy – Innovating for better saver outcomes.
  Available at: Digital, Data and Technology Strategy – Innovating for better saver outcomes

- **Financial Conduct Authority (FCA)** (2023).
  Artificial Intelligence in Financial Services.
  Available at: https://www.fca.org.uk

- **ISO/IEC** (2023).
  ISO/IEC 23894: Artificial Intelligence — Risk Management.
  International Organization for Standardization / International Electrotechnical Commission.

- **Royal Society** (2017).
  Machine Learning: The Power and Promise of Computers that Learn by Example.
  Available at: https://royalsociety.org

- Goodfellow, I., Bengio, Y. and Courville, A. (2016).
  Deep Learning.
  MIT Press.

- **Bender, E. M., Gebru, T., McMillan-Major, A. and Shmitchell, S.** (2021).
  On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?
  Proceedings of the ACM Conference on Fairness, Accountability and Transparency.

- **Russell, S. and Norvig, P.** (2021).
  Artificial Intelligence: A Modern Approach (4th ed.).
  Pearson.

- **Tzougas, G., John, A**., *et al.* (2022). *Classification of climate-related insurance claims using supervised learning*. Annals of the Institute and Faculty of Actuaries.

All references are non-proprietary, authoritative and believed by the authors to be appropriate to be included in a publication by members of the IFoA AIET Practice Board.

This document has been prepared by the Chair and the Council Representative of the AIET Practice Board. It represents the view of the authors and does not necessarily reflect the view of the IFoA, the Council of the IFoA, or the wider membership of the AIET Practice Board.