

ACTUARIAL INTELLIGENCE BULLETIN



January 2026

TABLE OF CONTENTS

| | |
|---|----|
| 2025 Toronto Seminar: “Emerging Issues & Practice for Actuaries in the Use of AI” | 2 |
| AI Innovations in Health Assessment for Life and Health Underwriting | 4 |
| AI Learning Agility: Practical Use Cases for Actuaries | 5 |
| SOA Research on the Emergence of AI in Actuarial Practice | 7 |
| The Society of Actuaries Elevates AI Education with PD Edge+ | 8 |
| IAA AI Task Force Update | 10 |
| From the SOA Research Institute | 13 |
| What AI Has Actually Changed in My Work..... | 15 |
| Further Insights from AI use in Actuarial Practice | 17 |
| What is Artificial General Intelligence (AGI) and What Will It Mean for the Insurance Industry? | 20 |
| Future Days: Points to Ponder about AI, Professionalism, and Actuarial Work..... | 21 |
| What’s in your AI? | 24 |
| GenAI Case Study: Leveraging Unstructured Claims Text with LLMs for Improved Cost Prediction..... | 26 |
| From Prototype to Production: A Practical Validation Playbook for GenAI in Actuarial Work..... | 28 |
| The Augmented Actuary in the Age of Agentic AI | 31 |
| Do you need AN explanation or THE explanation?..... | 32 |
| How Actuaries Are Engaging With AI: Insights from the SOA’s Summer 2025 Survey..... | 35 |
| Editors Award: Best Article of 2025..... | 38 |
| About the Society of Actuaries Research Institute | 40 |

Welcome to the January 2026 edition of the SOA Research Institute AI Bulletin! This bulletin serves as a platform for sharing knowledge and fostering collaboration around artificial intelligence within the actuarial community. Explore articles on strategic initiatives, practical tips, and research advancements, all aimed at empowering actuaries to leverage AI responsibly and effectively.

Caveat and Disclaimer

The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the Society of Actuaries Research Institute or the Society of Actuaries or its members. The Society of Actuaries Research Institute makes no representation or warranty to the accuracy of the information.

Copyright © 2026 by the Society of Actuaries Research Institute. All rights reserved.

2025 Toronto Seminar: “Emerging Issues & Practice for Actuaries in the Use of AI”

BEN MARSHALL, FSA, FCIA, CERA, MAAA

Start In early December, the SOA hosted a seminar about Actuaries and AI in Toronto. The event was an immediate sell out and the material presented was very well received, so we thought we would share some highlights from the meeting.

“Hiring” GenAI: What If We Governed AI Like a New Employee?

Bernice Lim, FSA, FCIA, CERA; Dean Rootenberg, FSA; and Ben Marshall, FSA, FCIA, CERA, MAAA

This presentation challenged actuaries to rethink model governance by comparing **Generative AI (GenAI)** to a human employee. The presenters suggested that while traditional Model Risk Management (MRM) focuses on inventory, rating, and validation for static tools, GenAI behaves more like a staff member that requires supervision, "hiring" criteria, and ongoing monitoring. The deck poses critical discussion questions, asking how AI models behave similarly to and different from employees and how governance strategies might shift in response, including the interaction between human workers and AI tools.

The session outlined a framework for discussing the tactical considerations of governing AI, asking attendees to evaluate where they stand in the adoption cycle and to identify "red flags" in AI adoption. The presentation emphasized that advancements in AI are blurring the lines between tools and "human-like cognition," necessitating a fundamental shift in how the industry approaches risk management.

Paper to Production – Lessons from the Field

Jonny Skerratt; Ryan Ali, FSA, FCIA, CERA; and Sarah Cheng, FSA, FCIA, MAAA

This session traced the journey of taking a GenAI-based tool from concept to production, covering technical and human factors. It highlights the "Data Dilemma," stressing that AI's quality is dependent on the data it learns from, amplifying existing quality issues. The process involves navigating the "Approval Gauntlet" of legal, ethical, and compliance checks to protect against bias and security risks. It also explored cultural readiness, stakeholder trust, and the need for new skills, emphasizing that AI reshapes roles and introduces new challenges regarding accountability for non-deterministic outputs. Companies are categorized by their AI adoption stage, from "Rejecting" at one end of the spectrum to "Optimizing" at the other.

AI Innovations in Health Assessment

June Quah, FSA, FCIA; and Jonathan Polon, FSA

See article following.

The Journey of Developing GenAI Products

Lan Tong, FSA, FCIA; Alan Wang, FSA, CERA, FCIA; and Green Chen, FSA, FCIA

This session explored the full lifecycle of developing GenAI applications in the actuarial and insurance domains. One speaker presented a "Lean In" mindset and partnering with AI experts. The other speaker emphasized "decentralized innovation" for productivity enhancement. Insights were shared on translating AI concepts into practical solutions that enhance user experience, modeling, governance, and decision-making, with a focus on ensuring actuarial soundness through validation.

The session also explored the **lifecycle of GenAI products**, addressing how to select use cases, design for user experience, and ensure actuarial soundness through validation. It highlighted the importance of governance frameworks that ensure trust and compliance, specifically asking how these processes differ from traditional modeling.

Not Real, but Useful: The Actuary's Guide to Synthetic Data

Harrison Jones, ASA; Bernice Lim, FSA, FCIA; and Tristan Walsh, ASA

This presentation introduced **synthetic data**—artificially generated data that mimics the statistical properties of real data—as a solution for data privacy and scarcity, enabling safer cross-border data sharing and system testing. The presenters discussed the trade-off between privacy and value, noting that while synthetic data improves privacy, optimizing utility remains a key challenge.

AI Learning Agility: Practical Use Cases for Actuaries

François Boulé, FSA, FCIA

See Article Following.

Generative AI for Business: Driving Growth and Competitive Advantage

Nazir Valani, FSA, FCIA, MAAA

This session was about how to take advantage of AI, the risks associated with AI, and how to create an AI business strategy to address business objectives. The session concluded with a strategic discussion on using Generative AI to create a "MOAT"—a defensive competitive advantage for the business.

SOA Research on the Emergence of AI in Actuarial Practice

Danielle Amiel, FSA, FCIA; Joe Alaimo, ASA, Hacia; JianGang He, FSA, FCIA; and Kevin Pledge, FSA, FIA

See Article following.

Accelerating Stochastic Calculations Through Neural Networks

Christopher Najjar, FSA, CERA; Martin Le Roux, FSA, FCIA, CFA; and Sally Xu, FSA, FCIA, CERA

This session covered two advanced applications of machine learning. The first section detailed using **Neural Networks (NN)** to price "Rainbow Options," demonstrating that NNs can significantly reduce runtime compared to Monte Carlo (MC) simulations while maintaining high accuracy.

The second section focused on modeling **Structured Asset Market Values** (like RMBS) using proxy functions calibrated via Least Squares Monte Carlo. It described a method using **Generalized Linear Models (GLM)** to estimate future market values at hundreds of projection dates without the computational cost of nested simulations. The results showed that these proxy functions provide accurate approximations of Monte Carlo valuations, with a highly efficient and robust calibration process.

The Evolving Actuary in the Age of AI

Dave Ingram, FSA, CERA

The final speaker looked ahead at a future for the profession, introducing the concept of the "**Actuary of the 5th Kind**" (**The Evolved Actuary**), who integrates human ingenuity, empathy, and wisdom with AI tools. Along with the "4th Kind" (AI Enhanced Actuary), who focuses on quality control and collaboration with AI. The presentation shared

four future scenarios, from an "AI Boom" driven by productivity gains to an "AI Depression" caused by job displacement and lack of retraining.

Ben Marshall, FSA, FCIA, CERA, MAAA is Regional Director Americas, Society of Actuaries



ACTUARIAL INTELLIGENCE BULLETIN

AI Innovations in Health Assessment for Life and Health Underwriting

Session Summary from the Toronto AI Seminar

JUNE QUAH, FSA, FCIA, JONATHAN P. POLON, FSA

Modern life and health underwriting relies on fast and inexpensive acquisition and analysis of high-signal health data to support real-time, accurate decisions.

Electronic Health Records

Electronic health records now represent one of the most expansive and information-rich data sources available to insurers. Traditional underwriting inputs—applications, paramedical exams, lab tests, MIB, and APSs—are increasingly supplemented or replaced by digital channels. Industry surveys show rapid growth in the adoption and evaluation of tools that classify or triage risk in accelerated underwriting, along with shifting views on which data sources provide the strongest protective value.

EHRs differ from legacy sources in both scale and complexity. They contain structured fields, semi-structured elements, and large volumes of unstructured clinical notes. While this breadth provides visibility into medical impairments, treatment patterns, and diagnostic histories, it also introduces operational challenges: extracting, standardizing, normalizing, and deduplicating data; determining which variables are materially predictive; and selecting appropriate analytical methods such as NLP, machine learning, and large language models.

The presenter illustrates how modern NLP and LLM systems transform narrative text into clean, structured outputs suitable for underwriting. These structured outputs can be consumed directly by underwriting rules engines, analytics layers, or automated decision pipelines, and can also generate concise summaries for human reviewers.

Because these capabilities introduce new operational and model-risk exposures, organizations must implement strong AI governance. Required elements include detailed documentation, model cards, clear definition of inputs and outputs, validation evidence, integration specifications, and explicit plans for change management, monitoring, and contingency operation. The overarching message is that EHR-driven underwriting is feasible and high-value, but only when paired with disciplined data engineering and controlled deployment of AI techniques.

Cough Analysis

Cough analysis is an inexpensive and non-invasive technology that can remotely screen people for specific respiratory diseases and conditions, such as tuberculosis, influenza, pneumonia or smoking, in less than a minute – based solely on the sound of their cough, submitted remotely via their own smartphone.

Cough analysis leverages the physiology of the human respiratory system. Irritation, inflammation, airway narrowing, and changes in mucus viscosity alter the frequency and harmonic structure of cough sounds. Each disease produces distinct impairments in respiratory mechanics, creating detectable acoustic patterns. These patterns can be visualized through waveforms, spectrograms, and model attention maps.

The screener is trained on a large, clinically verified dataset of 56,673 participants representing multiple continents, with illnesses confirmed through RT-PCR, chest imaging, antigen tests, and sputum cultures. Data preprocessing includes segmenting multi-cough samples, filtering noise, validating cough quality, and converting audio into Mel spectrograms and FFT images. Training incorporates staged dataset splits, and validation occurs through IRB-approved clinical studies where results are benchmarked against RT-PCR ground truth. Multiple clinical validation studies have been conducted, each typically enrolling 350–600 participants.

Applications for life and health insurers include:

- Underwriting: Real-time smoker detection and disease screening
- In-force Management: Periodic screening of policyholders for early detection of chronic respiratory conditions
- Group Benefits: Respiratory screenings for plan members and dependents
- Remote Patient Monitoring: Provide ongoing monitoring of respiratory function for people with chronic respiratory conditions

June Quah, FSA, FCIA, Vice President, Integrated Analytics at Munich Re; Jonathan Polon, FSA, Chief Analytics Officer at SigmaSight and Advisory Board Member at Raisonance



ACTUARIAL INTELLIGENCE BULLETIN

AI Learning Agility: Practical Use Cases for Actuaries

Session Summary from the Toronto AI Seminar

FRANÇOIS BOULÉ, FSA, FCIA

The central premise is that AI will not replace actuaries, but actuaries who do not leverage AI will be outperformed by augmented actuaries who pair professional judgment with AI-enabled efficiency.

Learning Agility: What It Is and Why It Matters

Learning agility is the capability to match the right AI tool to the right task, experiment thoughtfully, and adapt as technology evolves. Generative models excel at drafting, summarization, and structuring information; they are less reliable though for precise calculations, research validation, or statistics. The presentation stresses three pillars that build upon learning agility and are critical to optimal and proper AI use:

1. Learning Agility — Build a toolbox mindset; understand capabilities and limitations across tools.
2. Responsible Use — Work within vetted, secure environments; prioritize privacy, confidentiality, and compliance.
3. Human Judgment — Maintain peer review and actuarial judgment; verify, interpret, and own the results.

Use-Case Landscape

The session organizes practical AI applications into three categories, each grounded in typical actuarial tasks:

1) Technical Tasks

- **Coding & Debugging:** AI assists with code generation, optimization, documentation, and translation (e.g., Fortran into Python), improving speed, and reducing errors across coding languages. The session noted external predictions that AI could soon write most of the code, while current practice averages around 30% AI-generated code (as of November 2025), with adoption varying by team and task.
- **Dataset Creation & Claim Analysis:** AI helps extract, transform, and interpret structured and unstructured inputs, including free-text claim descriptions.
- **Research Workflows:** AI rapidly synthesizes multi-source material, reads charts/tables, and supports assumption setting and experience studies, always with robust peer review.

2) Communication & Documentation

- **Structuring Reports and Memos:** AI produces outlines, logical flows, and consistent formatting, accelerating drafting and reducing rework.
- **Text Generation & Rewriting:** From emails to technical notes, models adapt tone and style to audience needs while preserving clarity.
- **Translation at Scale:** A custom translation agent trained on 13,000+ actuarial sentences translated a 220-page report in under 2.5 minutes, delivering a reported BLEU score above 72, which surpasses typical human translators' baselines. Peer review remains essential to ensure context fidelity.
- **Presentation Creation:** A simple prompt or documents can be converted into professional slide decks, improving turnaround for stakeholder communications.

3) Other Support Functions

- **Exam Preparation:** AI accelerates study planning, practice questions generation and grading, and rapid concept review.
- **Image Generation:** Quick production of visuals for reports, infographics, and presentations.
- **Simultaneous Translation & Meeting Notes:** Real-time language support and automated note-taking to capture decisions, action items, and cross-team linkages, while enhancing collaboration and institutional memory.

Tooling Strategy and Evaluation

The presentation underscores that not all AI tools perform equally across tasks. Actuaries can benefit by comparing multiple solutions e.g., general platform vs specific tools. To help, a cheat sheet can be created that maps each use case to the tool that consistently delivered the best accuracy, speed, and ease of integration. Curated experimentation builds confidence while aligning with organizational guardrails.

Risks, Limitations, and Safeguards

AI can introduce material risks: bias, hallucinations, and inaccuracies. Confidentiality and data privacy are non-negotiable, especially for sensitive client or regulatory information. The session calls for:

- **Vetted Tools & Secure Environments** – Use enterprise-approved solutions; avoid ad-hoc uploading of protected content.
- **Professional Governance** – Maintain documentation of assumptions, versioning, and review logs.
- **Peer Review & Validation** – Treat AI outputs as starting points; verify sources, reconcile numbers, and confirm consistency.

The Augmented Actuary

The future belongs to actuaries who partner with AI. AI extends reach and speed; actuarial expertise ensures rigor, context, and accountability. The result is higher-quality analysis, clearer communication, and faster delivery, without compromising actuarial standards.

Key Takeaways for Practice

- High exposure = high opportunity: Actuarial tasks are especially amenable to AI augmentation.
- Start small, scale safely: Pilot vetted tools on low-risk tasks; build repeatable patterns and controls.
- Keep judgment central: Use AI to accelerate, not replace, critical thinking and peer review.
- Institutionalize learning agility: Provide training, shared playbooks, and a living tool-use map.

The bottom line is that AI is an ally. With learning agility, responsible use, and strong professional judgment, actuaries can convert AI's exposure into a durable advantage, becoming the augmented actuaries the profession needs.

Reference

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217, <https://doi.org/10.1002/smj.3286>.

François Boulé, FSA, FCIA, Office of the Chief Actuary



ACTUARIAL INTELLIGENCE BULLETIN

SOA Research on the Emergence of AI in Actuarial Practice

Session Summary from the Toronto AI Seminar

DANIELLE AMIEL, FCIA, FSA

Artificial intelligence (AI) is reshaping actuarial work and the profession's strategic direction. Moderated by Danielle Amiel (Medicus Pension Plan), the panel featured industry leaders Joe Alaimo (ProComp), JianGang He (Aon), and Kevin Pledge (Acceptiv), who shared insights on AI's impact on actuarial roles, standards, and organizational decision-making.

Understanding AI and Its Relevance

The discussion began with an overview of AI concepts, distinguishing traditional AI from Generative AI (GenAI) and Large Language Models (LLMs). While traditional AI focuses on predictive tasks like fraud detection and mortality modeling, GenAI introduces creative capabilities such as drafting reports and generating scenarios. These technologies are already influencing actuarial workflows, from underwriting to management reporting.

SOA Research Initiatives

The panel highlighted the SOA Research Institute’s strategic programs, which include Actuarial Innovation & Technology, Mortality & Longevity, Health Care Cost Trends, Catastrophe & Climate, Aging & Retirement, and Diversity, Equity & Inclusion. Within these programs, several AI-focused projects have emerged:

- Operationalizing GenAI for Actuaries – A guide for deploying and managing LLMs responsibly.
- AI Surveys – Baseline and longitudinal studies on AI adoption and sentiment among actuaries.
- Essay Collections – Exploring AI’s impact on actuarial practice, retirement, and longevity.
- Synthetic Medical Claims Data – Practical resources for creating synthetic datasets.
- Ethics for GenAI Models – Frameworks for fairness, transparency, and governance.
- AI in Healthcare and Insurance – Use cases and regulatory comparisons across global markets.

These initiatives aim to provide actuaries with tools, research, and ethical guidelines to navigate AI integration effectively.

Panel Insights and Audience Engagement

The dialogue addressed key questions:

- Benefits of AI for actuaries: Improved predictive accuracy, efficiency, risk insights, and cost reduction.
- Opportunities and challenges: GenAI offers new avenues for automation and personalization, but raises concerns about data privacy, bias, transparency, and over-reliance on automation.
- Skill evolution: Actuaries will need to develop competencies in data science, AI governance, and ethical risk management to remain relevant.
- Outlook: AI is expected to transform insurance distribution, underwriting, and reporting, while professional bodies and regulators must balance innovation with responsible governance.

Call to Action

The session concluded by encouraging actuaries to leverage SOA resources, participate in research, and pursue continuous learning to stay ahead in an AI-driven environment. Collaboration among organizations, regulators, and professional bodies will be critical to ensuring ethical and effective AI adoption.

Danielle Amiel, FCIA, FSA (she/her), Director, Governance & Risk Management, Medicus Pension Plan



**ACTUARIAL
INTELLIGENCE
BULLETIN**

The Society of Actuaries Elevates AI Education with PD Edge+

JON FORESTER, ASA, MAAA

Professional Development Edge+ (PD Edge+) is a learning product developed and delivered by the Society of Actuaries (SOA). Since its launch last year, it has become a reliable destination for actuaries seeking high-quality continuing education. Its steady growth reflects a broader shift in the profession: actuaries are increasingly engaging with artificial intelligence as an essential component of modern practice, and PDE+ is one of the places where that learning takes place.

AI as a Central Learning Theme

Across the platform, PDE+ has seen especially strong engagement with AI-focused material. Usage patterns indicate that actuaries are actively building fluency in emerging technologies, model governance, and practical AI applications.

The following list highlights the top 10 AI resources on PDE+ in 2025 (through October 2025.) These selections illustrate the breadth of content actuaries are leveraging to strengthen foundational knowledge and develop applied skills.

Top Ten AI Content on PD Edge+ in 2025

1. **Course: AI Dynamics, Evolution and Implementation**
Featuring [Jing Lang | LinkedIn](#)
2. **Short-form Video: Intro to AI**
Featuring [Jing Lang | LinkedIn](#)
3. **Course: AI Use Cases for Actuaries**
Featuring:
[Chris Smith, ASA, MAAA | LinkedIn](#)
[Han Henry Chen, FSA, FCIA, MAAA | LinkedIn](#)
[Joe Dorocak, ASA, MAAA | LinkedIn](#)
[R. Dale Hall, FSA, MAAA, CFA, CERA | LinkedIn](#)
4. **Short-form Video: Risk Identification**
Featuring [Dave Ingram | LinkedIn](#)
5. **Short-form Video: Generative AI Use Cases in Actuarial Science**
Featuring [Jing Lang | LinkedIn](#)
6. **AI Insights for Actuaries Webcast: The Essential Overview**
Developed from the SOA's *AI Insights for Actuaries* meeting
7. **Short-form Video: Mitigating AI Risks**
Featuring [Dave Ingram | LinkedIn](#)
8. **History of AI Recording**
Featuring [Jing Lang | LinkedIn](#)
9. **High-Level Techniques**
Featuring [Chris Smith, ASA, MAAA | LinkedIn](#)
10. **Model Validation**
Featuring [Chris Smith, ASA, MAAA | LinkedIn](#)

New and Noteworthy: Late 2025 Releases

PDE+ continues to expand with timely AI content that aligns with the evolving responsibilities of actuaries. These new courses address practical governance challenges, ethical considerations, and the role actuaries play in shaping responsible AI systems.

1. Navigating AI Bias

Featuring:

[Sherry Chan \(陳雪雯\), FSA, EA, MAAA, FCA | LinkedIn](#)

[Yukki Yeung | LinkedIn](#)

2. The Actuary in the Loop

Featuring, among others:

[Dave Ingram | LinkedIn](#)

[Amanda Hug | LinkedIn](#)

[Tom Callahan | LinkedIn](#)

[Robert Eaton | LinkedIn](#)

3. The Actuary in the Model

Featuring, among others:

[Dave Ingram | LinkedIn](#)

[R. Dale Hall, FSA, MAAA, CFA, CERA | LinkedIn](#)

[Arthur da Silva | LinkedIn](#)

[Igor Nikitin | LinkedIn](#)

The Path Forward

The SOA created PDE+ to support members pursuing ongoing professional development, and its expanding AI catalog reflects the profession's growing emphasis on emerging technology. As new courses and short-form videos are added, actuaries will have access to a curated set of resources that help them stay current on trends, strengthen their technical capabilities, and engage with AI responsibly. [Learn more](#) about PD Edge+.

Jon Forster, ASA, MAAA is Professional Development Learning & Content Design Director at the SOA.



ACTUARIAL INTELLIGENCE BULLETIN

IAA AI Task Force Update

FRANK CHANG, FCAS, PHD

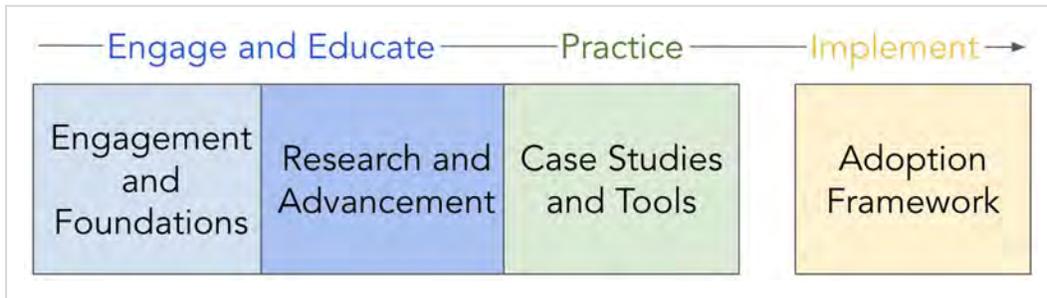
The IAA AI Task Force was formed subsequent to the [IAA Statement of Intent](#) with the following three objectives:¹

- **Advance the competency of the profession with respect to AI** by creating awareness of the risks and opportunities related to AI, facilitating knowledge sharing, and educating actuaries.
- **Promote the role of the actuary** in existing and emerging wider fields and raise its profile.
- **Prepare the IAA, as the voice of the global actuarial profession**, to proactively engage with Supranational organizations on AI-related risks and provide actuarial perspectives in their own related initiatives.

¹ International Actuarial Association, "Statement of Intent (SOI) for IAA Activities on Artificial Intelligence (AI)," approved December 14, 2023, PDF, https://actuaries.org/app/uploads/2025/04/SOI_AI.pdf, (Accessed January 8, 2026).

2025–2026 Workstreams

The task force is in its second year of work, organized into the following workstreams:



Engagement and Foundations

Workstream 1 (Engagement and Foundations) aims to create conversations about AI use in a professional practice context as well as assisting AI newcomers into the technology.

Our Engagement efforts are focused on two main projects. First, there is the development of a global online community where like-minded professionals can “Share, Discuss and Collaborate” on topics related to the use of AI by financial professionals. The website www.AlforActuaries.org is the home of our online community. Beyond supporting online discussion, AlforActuaries.org aims to help actuaries by highlighting top AI resource from actuarial and other organization around the world. Second, is the production of competitions where participants can learn-by-doing, using AI creatively in a friendly competition environment. After each competition, winners will share their techniques for the benefit of others.

We welcome all of you to join us at aiforactuarie.org.

Our Foundations track is working to develop learning paths for the different types of AI and their related skills. These learning paths will provide curated learning materials to help those new to AI. A challenge for those new to AI is first navigating the “what is what” with the many flavors of AI and then finding good resources to learn about these technologies. The Foundations learning paths has the goal of simplifying the process of building a strong foundational understanding of the branches of AI relevant to them.

Research and Advancement

Workstream 2 (Research and Advancement) divides its deliberations into three areas: the curation of AI-related knowledge & creation of an AI-related knowledge base example; the definition of relevant research areas for the AI-enabled actuary; and the creation of a sustainable communication framework regarding research and advancement in AI.

The AI-related knowledge base example, which has an international focus, is currently being prototyped. Pending approval, it is also planned to include an AI agent use case to optimize usage and accessibility.

Several high-end skill and competency categories have been discussed and collected in the second sub-group to support the AI-enabled actuary for example (automation of routine tasks, changing interpretation of actuarial work, obtaining buy-in from external stakeholders, ethical considerations etc.). Close cooperation with other workstreams ensures relevant research is aligned and accessible.

AI is arguably the most important topic for the actuarial profession right now. In line with this perception, a great deal of development and research is being carried out by individuals, actuarial associations and companies around

the world. The subgroup aims to propose an initial (for example direct communication to FMAs) framework for ongoing collaboration on research-related activities in the AI actuarial space.

Case Studies and Tools

Workstream 3 (Case Studies and Tools) focuses on developing practical resources that help actuaries understand and apply AI in their daily work. Its activities center on three areas: the development of actuarial AI case studies, the curation and structuring of existing materials, and the development of guidance on AI tools relevant to the actuarial profession.

The Workstream identifies gaps in the current landscape and produces case studies that address real actuarial needs across practice areas, ensuring they are accessible, reproducible, and aligned with broader IAA AI Task Force objectives.

A complementary focus is the identification, curation, and development of educational material on AI tools – such as conversational chatbots (e.g., ChatGPT) and coding assistants (e.g., Cursor) — to support actuaries in adopting these technologies responsibly and effectively.

The Workstream also maintains the [IAA AI Task Force's GitHub repositories](#) on case studies and tools, providing templates (e.g., Jupiter Notebook, R Markdown) and guidance to enable contributions from the global actuarial community. In the September issue of the SOA AI Bulletin, Workstream 3 introduced and explained this GitHub resource, and another article in the current issue presents one of our newly developed case studies in more detail.

Adoption Framework

The Adoption Framework Workstream (Workstream 4) plays a critical role in equipping actuaries globally with a structured methodology to evaluate AI adoption strategies and integrate responsible AI practices. The primary goal is to develop an Interactive Framework that serves as a comprehensive end-to-end guide. This guide will help actuaries synthesize essential considerations - including best practices, professionalism, governance, ethics, and practicality.

By operationalizing previous outputs, the workstream aims to expand the actuary's role into wider fields, raise the profession's profile, and support the International Actuarial Association (IAA) in engaging regulators and supranational organizations on AI-related risks. The team is currently integrating existing deliverables (such as papers on governance, testing, documentation, and professional considerations) into this framework. The design covers the entire AI implementation lifecycle, from Model Strategy & Planning and Data Management to Model Monitoring and Maintenance & Retirement.

Key areas of focus also include strengthening regulatory alignment, updating guidance on regulatory matters, and addressing emerging topics in the field of narrow and generative AI, as well as defining disclosure practices for transparency and explainability.

Frank Chang, FCAS, PhD is Vice President, Applied Science, Uber.



**ACTUARIAL
INTELLIGENCE
BULLETIN**

From the SOA Research Institute

[2025 AI: A Collection of Essays²](#)

A collection of essays exploring AI's real-world impact on actuarial work.

First Place Prize Winners (Tied)

[The Intern's Intern: AI's Role in Developing Early-Career Actuaries](#)

Nii Amoo Decardi Nelson

AI functions best as a software engineer specialist rather than a generalist, accelerating output but demanding that early-career actuaries shift their focus from manual execution to rigorous auditing, ethical validation, and governing the AI's complex outputs.

[Experience Studies Harnessing an AI Agent — A Proof-of-Concept Lightyears Past Code Generation](#) AI can act as a nearly autonomous agent, directly generating and executing code in a Python notebook to perform complex experience study analysis, moving beyond its traditional role as a simple code-generating tool.

Second Place Prize Winner

[The Actuary and the Algorithm: Navigating the New Symbiosis of Judgment](#)

Niranjan Rajandran

AI is transforming the actuary's role from a technician to a translator and conductor, automating routine work but elevating the need for human judgment to strategically interpret data, explain model opacity, and build governance frameworks for complex algorithms.

Remaining Essays

[AI in Practice: Building Practical Solutions for a Resource-Strapped](#)

Shaun Crossman, FIA, FASSA, CERA, FRM

For a smaller insurer, AI provides practical, incremental wins by automating tasks involving messy data and enabling complex projects like lifetime value modeling that were previously impractical with traditional tools, but its success requires persistence, human oversight, and adherence to governance.

[AI Assistants and Simulations](#)

Dave Ingram, FSA, MAAA, CERA

AI assistants, when customized with prompts and data files, offer a powerful, repeatable workflow for actuaries beyond simple search, notably enabling engaging business simulations for practicing soft skills and critical thinking.

[Bridging the Gap: How AI Changed My View of Actuarial Work](#)

Sathiya Livingston, FSA, MAAA, CERA, PMP

From an operational perspective, AI's introduction has shifted the actuarial role from one of technical guardians to one of trust builders and explainers, focused on questioning opaque models and ensuring fairness and accountability to clients and regulators. Further reflections following.

[Insights from AI Use in Actuarial Practice](#)

Prabhdeep Singh, FSA, MAAA, CERA, PMP

² Society of Actuaries Research Institute, November 2025, *2025 AI: A Collection of Essays*, <https://www.soa.org/resources/research-reports/2025/ai-insights-actuarial-practice/>, (accessed January 8, 2026).

AI serves as a powerful accelerator for coding actuarial functions and learning new technical tools, but its confident errors and lack of design judgment necessitate that actuaries apply rigorous professional skepticism and validation. Further reflections following.

[Actuarial A.I 1.0](#)

Nathan Worrell, FSA; Green Chen, FSA; Brandon Lin, FSA

Creating an internal AI "Navigator" with restricted documentation can successfully serve as a learning and documentation aid, despite challenges with actuarial jargon and incomplete knowledge, by efficiently compiling information from disparate sources and ultimately improving customer support.

[SOA Member AI Survey — Summer 2025³](#)

A survey of SOA members exploring AI adoption, use, and perception across the actuarial profession.

[Provider Use of AI in Healthcare⁴](#)

This report examines how healthcare providers are using AI—including predictive, generative, and agentic systems—to improve outcomes and efficiency and highlights the evolving role of actuaries.

[Healthcare AI—Current Applications and What’s Next: An Expert Panel Discussion⁵](#)

Expert panel discusses AI adoption in healthcare, focusing on clinical impact and governance.

[Insights from an Expert Panel on AI and Actuarial Responsibility⁶](#)

A panel-driven exploration of AI bias, equity, and the evolving responsibilities of actuaries in healthcare and insurance.

[Americas AI Action Plan: AI Insights and Insurtech Insights⁷](#)

The implications of the White House’s July 2025 AI Action Plan, *Winning the Race — America’s AI Action Plan*, and what it means for actuaries, insurers, and the broader financial sector.



ACTUARIAL INTELLIGENCE BULLETIN

³ Joe Alaimo, ASA, ACIA, November 2025. *AI Survey: Summer 2025*, Society of Actuaries Research Institute, <https://www.soa.org/resources/research-reports/2025/ai-member-survey-summer-2025/> (accessed January 8, 2026).

⁴ Joe Dorocak and Don McLellan, October 2025, *Provider Use of AI in Healthcare*, Society of Actuaries Research Institute, <https://www.soa.org/resources/research-reports/2025/provider-use-ai-healthcare/>, (accessed January 8, 2026).

⁵ Ronald Poon Affat, FSA, FIA, MAAA, CFA, HIBA, October 2025, *Healthcare AI—Current Applications and What’s Next: An Expert Panel Discussion*, Society of Actuaries Research Institute, <https://www.soa.org/resources/research-reports/2025/healthcare-ai-expert-panel/>, (accessed January 8, 2026).

⁶ Ronald Poon Affat, FSA, FIA, MAAA, October 2025, *From Health Inequities to Societal Bias: Insights from an Expert Panel on AI and Actuarial Responsibility*, Society of Actuaries Research Institute, <https://www.soa.org/resources/research-reports/2025/ai-actuarial-bias-equity/>, (accessed January 8, 2026).

⁷ Frances M. Green and Dale Hall, November 12, 2025, “Winning the Race — America’s AI Action Plan,” Society of Actuaries Research Institute, <https://getpluggedin.libsyn.com/winning-the-race-americas-ai-action-plan>, (accessed January 8, 2026).

What AI Has Actually Changed in My Work

Further reflections to the recently published Essay

SATHIYA LIVINGSTON FSA, MAAA, CERA, PMP

Months after writing the essay, I’ve realised that the real changes weren’t about efficiency at all. The surprising shift has been in how I approach the work I share with actuaries—and how our collaborations have evolved because of AI. The tasks may look the same, but the entry point into those tasks feels very different now.

The End of Starting from Zero

One of the simplest but most meaningful changes AI brought into my day-to-day work is the disappearance of the blank page. Whether it’s preparing a client note, drafting a requirements document, or outlining a model review discussion, I rarely start empty anymore. There’s always a rough draft or structured starting point.

That alone changes the tone of the work. I’m not stuck thinking, *where do I begin?* Instead, I’m already in problem-solving mode:

What’s missing? What’s inaccurate? What matters most for the actuary or client on the other side of this conversation?

For someone who interacts across underwriting, claims, operations, and actuarial teams, that shift in momentum has been surprisingly empowering.

Speed Changes Expectations — Even When No One Says So

This is where I’ve seen actuaries shine: asking better questions, catching subtle inconsistencies, drawing lines between fairness and accuracy, and reframing uncertainty for clients. AI hasn’t replaced their judgment—it’s made the value of that judgment more visible.

AI speeds up routine work, but the deeper shift is psychological. Once you realise a task that took half a day now takes an hour, you start raising the bar for yourself—even if nobody else asks you to.

I now spend more time refining rather than rushing.

More time clarifying assumptions.

More time ensuring the actuarial team has what they need upfront, not after three rounds of email.

But speed also brings a new kind of responsibility: AI can produce several versions of a document, but **I still have to choose the one I will stand behind** when talking to an actuary, a client, or a stakeholder.

AI makes the draft.

Humans make the decision.

And the decision still carries the weight.

Curation: The New Hidden Skill

I didn’t expect the biggest change in my work with actuaries to be *curation*, but that’s where I’ve landed.

AI can generate explanations, talking points, classification logic, scenario outlines, and documentation. But quality? Relevance? Accuracy? Those still depend on human filtering.

My job increasingly feels like selecting, shaping, and refining rather than creating from scratch. I find myself asking:

- Is this explanation technically, right?
- Would an actuary accept this reasoning?
- Is this output misleadingly confident?
- What matters for the business problem we're trying to solve?

Curation wasn't in my job description, but it has quietly become one of the most important parts of supporting actuarial work.

Where Human Judgment Shows Up More Clearly

Using AI regularly has made me appreciate the interpretive, human parts of the work far more—especially when collaborating with actuaries.

Some examples stand out:

- Explaining a surprising model score to a client who expects clear reasons.
- Helping translate operational realities into actuarial assumptions.
- Spotting when AI misclassifies a risk because it lacks context.
- Reconciling mismatches between the data and what we know from field experience.

AI does not solve these moments. It only sets the stage for them.

This is where I've seen actuaries shine: asking better questions, catching subtle inconsistencies, drawing lines between fairness and accuracy, and reframing uncertainty for clients. AI hasn't replaced their judgment—it's made the value of that judgment more visible.

Collaboration Looks Completely Different Now

One of the biggest surprises has been how much AI changed the feeling of collaboration.

When I sit down with an actuary now, we're often starting further ahead. Whether the starting material came from me or them, we usually have:

- a summary,
- a comparison table,
- a draft calculation flow,
- or a question-driven outline.

That means we enter the conversation already aligned on context. Instead of catching up, the discussion jumps straight to interpretation, prioritisation, and decision-making.

Junior analysts—on both the operations and actuarial sides—particularly benefit. AI doesn't give them experience, but it gives them enough structure to join the conversation with confidence instead of hesitation.

- AI doesn't flatten expertise.
- But it flattens the starting line.
- And that's healthy for cross-functional work.

Curiosity Is Suddenly a Real Advantage

AI has made it dramatically easier to explore unfamiliar topics. If I come across a regulatory change, a new modelling concept, or a pricing philosophy I've never worked with, I don't need to wait for a deep-dive call or search through dense documentation.

A few minutes of targeted prompting gives me a workable map.

That means I show up to actuarial discussions more prepared, with better questions, and with a clearer sense of what's at stake. Curiosity no longer feels like a bonus trait—it has become a practical edge in keeping up with an evolving field.

What I'm Certain About Now

In my first essay, I wrote that AI was reshaping actuarial work. After several months, I'd revise that thought:

AI didn't reshape actuarial work. It revealed what actuarial work is truly built on.

- Sound reasoning.
- Good judgment.
- Ethical responsibility.
- Clear communication.

AI made the scaffolding faster—but the core stayed the same.

And from my side of the table, supporting actuarial teams and helping bridge operations and analytics, that core has never been more visible.

A Final Thought

If the first wave of AI was about speeding up tasks, the next wave will be about sharpening roles — including mine, as someone who works closely with actuaries but isn't one myself.

- The repetitive layers are fading.
- The interpretive layers are rising.
- And the collaborations between actuarial, operations, underwriting, and analytics feel more meaningful, not less.

AI isn't just helping us calculate the future. It's helping us understand our place in shaping it — together.

Sathiya Livingston FSA, MAAA, CERA, PMP is Business Development Manager at Flatworld Solutions Pvt Ltd



**ACTUARIAL
INTELLIGENCE
BULLETIN**

Further Insights from AI use in Actuarial Practice

Reflections to the recently published essay

PRABHDEEP SINGH, FSA, CERA, MAAA, PMP

My early workflow was straightforward: break programming tasks into small, clear instructions and expect the AI to execute them. It significantly improved my productivity, but it pulled me down to the mechanics of Python rather than focus on the structure of actuarial models.

The breakthrough came when I learned to ask the AI to review my system specifications and not just to implement them. I now draft purpose-built specifications for entire modules, and before writing a single line of code, I ask the assistant to identify ambiguities, contradictions, or missing information from my specs.

This forced me into a level of design rigor I had previously avoided. I no longer had to think through Python implementation details while trying to think like an actuary. The AI took on the programming domain; I took on the modeling domain. The result is an architecture that reflects my intent, clarified and pressure-tested by the assistant's questions.

This was my first glimpse of what a *symbiotic relationship* with AI might look like.

An MVP Pricing Model in Weeks: The Turning Point

The best illustration of this workflow came when I built a full pricing model for a single-premium whole life product. The model spans multiple modules:

- Assumptions
- Model office
- Cash values
- Reserves
- Product design and illustration
- Liability cash flows
- A projection engine
- A solver for target profit metrics and product performance
- An output module producing Excel-based projected financial statements

Building this system would previously have required months. With my AI agent, it took a couple of *weeks*—a genuine turning point. The emotion that best describes this moment is **empowerment**. The result was strong enough that I shared a recorded AI workflow with a prospect, supported by a written concept paper I had prepared.

This experience sharpened my long-term ambition: to help expand the actuarial profession's modeling capabilities by developing AI-ready actuarial frameworks.

AI as a Design Collaborator—Sometimes Brilliant, Sometimes Stuck

As I moved from individual functions to system-level design, one theme became clear: AI can be unexpectedly insightful.

When I explained that the goal was an MVP for a prospect, the assistant suggested skipping a full QuantLib asset model and instead assuming a simple return on a cash account. It was a simplification I initially resisted but eventually recognized as appropriate for a prototype.

But the inconsistency remains. In one case, while unit-testing a commutation function, it got stuck in a loop of modifying the test or the code. I decided to stop the loop and just manually adjust the unit test.

Working effectively with AI requires understanding this: it can act like a brilliant actuarial student at times and remind you that it is just an automation at other times.

The Biggest Lesson: The Less AI Guesses the Better It Gets

One of the most important insights I gained is this: to use AI effectively, I must explicitly ask it for **feedback on my instructions or specifications**.

Humans provide feedback implicitly— e.g. through facial expressions. AI does not unless asked to do so by the actuary. Without instruction, it will *guess* what it should do when faced with an ambiguity. Then, the actuary will more likely not get what he expected from AI and get frustrated.

Reframing AI: Not a Tool, but a Colleague with Limits

The biggest misconception I see among actuaries is viewing AI as a productivity tool akin to spreadsheet programs replacing calculators or e-mail replacing typed memos. GenAI is different because of its adaptive and probabilistic nature. It adapts to how the actuary works and expands not just what he can do but also facilitates the expansion of his knowledge and skillset. If done right, it does start feeling like a colleague – an abstraction that is not possible with tools. If the AI agent has not started feeling like a colleague in the actuary’s work, then the actuary has not started fully leveraging it.

The actuary sets intent, architecture, and judgment; the AI expands capacity, speed, and reach.

Closing Reflection

If the first essay was about courage and experimentation, this follow-up is about capability and understanding. In just two months, my work with AI has matured from scattered trials to structured modeling. I learned that developing actuarial systems with AI is not merely a question of prompting—it is a question of *collaboration*, where both partners have distinct roles.

The next phase for the profession is clear: actuaries must learn to work with AI agents in a way that respects both their strengths and their limitations. And the only path to that understanding is through experience—trying, refining, and learning how to build systems together.

The actuary who learns to design with AI, rather than around it, will not just be more productive. They will help shape the next era of actuarial modeling.

Disclaimer: At least some parts of this essay were produced with the assistance of GenAI.

Prabhdeep Singh, FSA, CERA, MAAA, PMP Consulting Actuary, Singh Actuarial Consulting, LLC.



**ACTUARIAL
INTELLIGENCE
BULLETIN**

What is Artificial General Intelligence (AGI) and What Will It Mean for the Insurance Industry?

JEFF HEATON, PH.D

Just over three years ago, on November 30, 2022, ChatGPT was released, and for many, the era of AI began. The insurance industry has seen what the domain of actuaries, data scientists, and software engineers was once brought to everyone, even associates who were not traditionally “technical.” This democratization of advanced analytics and automation has reshaped not only how work is performed but who can participate. The speed of change and innovation seems breakneck, and we are seeing new terminology describing future states of even more powerful AI—Artificial Generative Intelligence (AGI), the next step on the road to superintelligent machines.

AGI promises to have profound impacts on insurance. A new era is on the horizon; preparing for it is a shared responsibility.

In late 2025, AI is everywhere in our lives and quickly becoming just another part of how we work and make decisions each day. When we drive to the office, our car likely assists or even drives on its own. Meanwhile, we might receive several instant messages from family and colleagues, which our mobile device’s AI assistant conveniently categorizes and prioritizes. Once we reach the office, AI is there to help. For example, an underwriter might use an AI tool to analyze an EKG in an underwriting file and use an email assistant to draft a response for follow-up medical records.

This experience may feel very uniform; however, it is made possible by multiple AI models stitched together using traditional computer programming. This stitching is a big part of agentic AI, which allows users to connect very specialized AI models into a single cohesive solution. Agentic AI also allows these models to reach out into the “real world” to carry out tasks, the very definitions of “agency.”

Modern AI is made up of several models that are essentially “one-trick ponies.” Stable diffusion generates all the fascinating images, for example, while other generative models specialize in video synthesis, text generation, speech recognition, and even code generation. Separate systems power real-time language translation, self-driving cars, and medical image interpretation. Each excels at a narrow task, but when combined through the “stitching” of agentic AI, they form the broader digital ecosystem we now take for granted.

These models are specialists that can currently only communicate at the most superficial level. A very good example of this is chess playing. Machines dominate chess; there is simply no hope of a human ever beating even a chess computer running on a mobile device at its top skill level. The book *Game Changer* walks readers through the unconventional strategies discovered by AlphaZero, many of which were so unusual that human grandmasters had never considered them.

But here’s the catch: that book was written by Matthew Sadler and Natasha Regan, both human beings. The most advanced AI chess engines can’t actually write, and we can’t just download their strategies. Writing a book is the domain of LLMs, such as ChatGPT. For its part, ChatGPT plays chess surprisingly poorly and has even been beaten by versions of chess running on very old hardware, such as the Atari 2600. To write that book, the authors had to observe AI playing hundreds of games, much like Jane Goodall learning gorilla behavior through observation in the wild.

Future models will consolidate these functions, much as your mobile device has become your GPS, camera, phone, pager, note taker, personal recorder, and laptop replacement. This will result in a new model capable of performing tasks without the need for agentic code to stitch multiple specialist models together. If this future model type is

truly general, it will be immediately applicable to new problem domains we have yet to conceive. This may include adapting to the changing regulatory and market conditions of the insurance industry to advance automated underwriting, actuarial projection, and even statutory accounting. We will have entered the era of Artificial General Intelligence (AGI).

Will AGI be as smart or smarter than humans? That is unlikely; it could well be smarter than current AI, just much more general. Surpassing human intelligence will require another breakthrough called recursive self-improvement (RSI), in which AI completely masters computer programming, as it did with chess, and no longer requires human help to improve. RSI is often called “the singularity” and will quickly lead to an AI that is superintelligent – beyond human intelligence.

It may take some time before we achieve AGI, much less RSI or superintelligence. Estimates for the arrival of AGI range from a few years to a few decades. AGI will be a breakthrough, and technological breakthroughs do not occur on a fixed schedule. On January 31, 2015, a *Forbes* article predicted that “The (truck) driver is practically no longer required,” yet in 2025, according to the Bureau of Labor Statistics, employment of heavy and tractor-trailer truck drivers is projected to grow 4% from 2024 to 2034. We see similar predictions for many professions in 2025.

So, when will AGI happen? Crowdsourced forecasting platforms can be fairly accurate at these sorts of predictions, as they poll the public and give higher weights to participants with higher past accuracy. One such leading platform, “Metaculus Prediction,” gave a median date of May 2033; the question originated in 2020, and this date will fluctuate as we move forward. Once AGI arrives, our systems will become highly adaptable, and our world will be forever changed.

As with virtually every other industry, AGI promises to have profound impacts on insurance. Given the rapidly evolving nature of bioinformatics and regulatory practices, insurers will be well served to closely monitor this technology’s progress. A new era is on the horizon; preparing for it is a shared responsibility.

Jeff Heaton, Ph.D., VP AI Innovation, RGA



Future Days: Points to Ponder about AI, Professionalism, and Actuarial Work

MITCH STEPHENSON, FSA, MAAA

I have observed a diversity of sentiments about AI. These range from “overhyped” to “necessary.” [According to the World Economic Forum](#), we may be at the beginning of the fourth industrial revolution.⁸ The previous three were factories (late 1700’s – mid 1800’s), assembly lines (late 1800’s – early 1900’s), and personal computing (mid 1900’s

⁸ Klaus Schwab, “The Fourth Industrial Revolution: What It Means, How to Respond,” *World Economic Forum*, January 14, 2016, <https://www.weforum.org/stories/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/>, (accessed January 8, 2026).

– early 2000’s). Given the profound impact AI will have on our personal and professional lives, here are background and key points that audiences often leave me to ponder.

AI Ethical Principles

The ethical considerations of using AI are not new. According to [The Medium](#), ethical implications of using AI included automation and job displacement in the 1970-1990’s.⁹ As paper records moved to electronic forums in the 1990’s-2000’s, ethical considerations included privacy and data ownership. The appropriate use of facial recognition became a concern in the 2010’s. Autonomous vehicle accidents and the use of personal data in the Cambridge Analytica scandal in 2018 generated ethical considerations. Since then, a significant industry trend has been for companies, agencies, and organizations to release their own ethical principles to govern the responsible use of AI.

Audiences to whom I present often articulate the significance of avoiding bias and protecting privacy as being chief among their concerns. According to [my research](#), there are three other commonly articulated ethical principles I observe.¹⁰ One is to ensure there is transparency. This is necessary both in documentation of AI tools and in ensuring customers know they are interacting with AI. Another is to make sure tools are reliable and safe. This includes ensuring there is ongoing performance monitoring and available backups. The last is to ensure there is accountability. This includes ensuring there is a human in the loop where appropriate. Occasionally, audience members articulate additional principles which should drive how we develop, deploy, and use AI.

Point to ponder #1: Which additional ethical principles would you articulate as they relate to the use of AI in actuarial work?

Actuarial Professionalism Standards

Actuarial audiences, presenters, and authors typically identify four Actuarial Standards of Practice (ASOPs) as being most relevant to the use of AI. These are: ASOP 12 (Risk Classification), ASOP 23 (Data Quality), ASOP 41 (Actuarial Communication), and ASOP 56 (Modeling). Code of Conduct precepts commonly cited as relevant to AI work include the preamble, Precept 1 (Professional Integrity), Precept 2 (Qualifications), Precept 3 (Actuarial Standards), Precept 4 (Communication & Disclosures), Precept 8 (Prevent Misuse of Work), & Precept 9 (Confidentiality). Of the above, the Actuarial Standards Board (ASB) released or revised only ASOP 56 after 2018, before the prevalence of AI ethical principles in the industry.

Even so, the Code of Conduct and ASOPs address each of the five ethical principles articulated above. The natural cycle of ASOP revisions — [driven by the ASB](#)¹¹ — which “reviews and evaluates current and emerging practices” — will allow for updates to ASOPs over time. Recent developments around Generative AI create a challenge of speed, comprehensiveness, and relevance for this updating process. Fortunately, there is a large body of relevant emerging work made available by professional organizations. This includes [three papers by the International Association of Actuaries](#) announced in December 2025.¹² The American Academy of Actuaries (AAA) released [Actuarial](#)

⁹ Kai Kaushik, “Tracing the Evolution of AI Ethics Through Time,” Medium, January 16, 2024, <https://medium.com/@kumarakaushik/tracing-the-evolution-of-ai-ethics-through-time-b7464d7dcd55>, (accessed January 8, 2026).

¹⁰ Mitchell Stephenson, “Zero to AI Governance: Establishing a Principles-Based Framework,” in *AI Bulletin* (May 2025) (Society of Actuaries Research Institute, 2025), 7–9, PDF, <https://www.soa.org/49955c/globalassets/assets/files/resources/research-report/2025/2025-05-ai-bulletin.pdf>, (accessed January 8, 2026).

¹¹ Actuarial Standards Board, “About ASB,” <https://www.actuarialstandardsboard.org/about-asb/>. ([actuarialstandardsboard.org](https://www.actuarialstandardsboard.org/)), (accessed January 8, 2026).

¹² International Actuarial Association, “IAA Releases Three New Papers to Support Responsible AI in Actuarial Practice,” news release, December 3, 2025, <https://actuaries.org/news-post/iaa-releases-three-new-papers-to-support-responsible-ai-in-actuarial-practice/>, (accessed January 8, 2026).

[Professionalism Considerations for Generative AI](#) in 2024.¹³ The Society of Actuaries Research Institute maintains a collection of [Artificial Intelligence Research Reports](#), most recently updated in November 2025.¹⁴

Point to ponder #2: Does the actuarial profession need an AI-specific ASOP, a near term update of existing relevant ASOPs, or to rely on professionalism standards “as is” while referencing other relevant resources?

Other Professionalism Considerations

At a recent presentation I attended, participants indicated they use AI in their work primarily for research, coding, and summarization. The authors of the AAA Actuarial Professionalism Considerations for Generative AI paper address considerations for each of these uses. One consideration is: *If you rely on AI tools for your actuarial findings, how did you determine that reliance is appropriate? Could that reliance withstand a professional, regulatory, or audit challenge?* There is an often-cited [2023 case](#) of lawyers submitting a legal brief composed by ChatGPT.¹⁵ The tool invented fictitious case citations to support their argument. The lawyers were subsequently sanctioned. More recently, per [US News](#), a study found that 490 law briefs submitted in 2025 contained hallucinations.¹⁶ This means the AI responses “contain false or misleading information.”

It is important that any actuary relying on Generative AI for research, summarization, or coding can stand behind the content therein. Another consideration articulated in the AAA paper is: *Is using a GenAI model appropriate for the assignment, given its costs and limitations? Would a simpler model suffice?* This may be increasingly relevant as actuaries become more comfortable using AI coding tools like GitHub Copilot, Python, and SageMaker. According to the [AI Journal](#), AI has coded an entire video game, written full books, and created new languages. Its coding capabilities are compelling in terms of speed, efficiency, and capabilities for model development.

Point to ponder #3: Under what circumstance would you use AI as a primary coding device for converting a block of business from old software or coding a new product? What safeguards would you put in place, and when would this use of AI not be appropriate?

AI as an Opportunity

According to a [study by McKinsey](#), 92% of companies plan to increase their AI investments over the next three years.¹⁷ An Infosys study cited in the [Economic Times](#) indicates only 2% of 1,500 respondent companies meet Responsible Artificial Intelligence standards.¹⁸

Mitchell Stephenson, FSA, MAAA is Vice President, Head of Model Governance Solutions at Fannie Mae Name

¹³ American Academy of Actuaries, Committee on Professional Responsibility, *Actuarial Professionalism Considerations for Generative AI: A Professionalism Discussion Paper* (September 2024), PDF, <https://actuary.org/wp-content/uploads/2024/10/professionalism-paper-generative-ai.pdf>, (accessed January 8, 2026).

¹⁴ Society of Actuaries Research Institute, “Artificial Intelligence Research Reports,” Society of Actuaries, <https://www.soa.org/research/topics/artificial-intelligence-res-report-list/>, (accessed January 8, 2026).

¹⁵ Sara Merken, “New York Lawyers Sanctioned for Using Fake ChatGPT Cases in Legal Brief,” Reuters, June 26, 2023, <https://www.reuters.com/legal/new-york-lawyers-sanctioned-using-fake-chatgpt-cases-legal-brief-2023-06-22/>, (accessed January 8, 2026).

¹⁶ Cathy Bussewitz, “Mistake-Filled Legal Briefs Show the Limits of Relying on AI Tools at Work,” *U.S. News & World Report*, October 30, 2025, <https://www.usnews.com/news/technology/articles/2025-10-30/mistake-filled-legal-briefs-show-the-limits-of-relying-on-ai-tools-at-work>, (accessed January 8, 2026).

¹⁷ Hannah Mayer, Lareina Yee, Michael Chui, and Roger Roberts, *Superagency in the Workplace: Empowering People to Unlock AI’s Full Potential* (McKinsey & Company, January 28, 2025), accessed January 8, 2026, <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work>, (accessed January 8, 2026).

¹⁸ ETtech, “AI Mishaps Hit 95% Executives, Only 2% Firms Meet Responsible Use Standards: Infosys Study,” *The Economic Times*, August 14, 2025, <https://economictimes.indiatimes.com/tech/artificial-intelligence/ai-mishaps-hit-95-executives-only-2-firms-meet-responsible-use-standards-infosys-study/articleshow/123305693.cms>, (accessed January 8, 2026).

ACTUARIAL INTELLIGENCE BULLETIN

What's in your AI?

RAYMOND K. SHEH PHD, FRANCES M. GREEN ESQ., LL.M., AND KAREN GEAPPEN MCSSD

When you buy a box of cookies, the label tells you what's inside: flour, sugar, butter, and eggs. There may be a "nut-free" assurance. If contamination is found, sources can be traced, and batches recalled. When an AI system generates a summary of an important email or flags a fraudulent transaction, what "ingredients" went into that output? If something goes wrong, can the source be traced? As AI systems are increasingly integrated, explicitly or implicitly, into decision-making, we face a growing challenge in assessing, weighing, and managing the risks they entail.

The Challenges of Hidden Complexity

An actuary might use an AI system to draft regulatory reports, transforming tables of reserve calculations and assumption changes into clear explanations for state insurance departments. On a more sophisticated level, an actuary could deploy an AI agent to continuously monitor emerging mortality data across multiple databases, automatically flagging significant deviations from expected trends and assembling preliminary impact analyses on life insurance reserves. How can the risks associated with something important being omitted or misrepresented be managed?

AI systems have an intricate, dynamic web of dependencies beyond those for traditional software. The "ingredients" contributing to each output may span dozens of entities across disparate industries and contexts. Even if the AI system was trained on high-quality data sources such as historic reports and regulatory texts, the foundational models that enable it to understand language may be trained on web-scale datasets that include historically biased data, jokes, sarcasm, humor, and incorrect homework answers posted to public forums.

Beyond data quality and ethics concerns, these web-scale datasets are often just uncurated links to data hosts, such as websites and social media platforms, making them susceptible to "data poisoning" attacks. These include attackers registering expired domains or social media accounts and replacing previously valid data with their own before AI systems retrain on that dataset.

Hidden Complexity Equals Uncertain Liability

This diffusion of liability, borne of inherent technical opacity, stymies the very idea of readily regulating or traditionally "managing" supply chain activity and ultimately assessing liability to enjoin or punish the bad or negligent actor(s) who have caused harm or endangered safe use. Traditional supply chain liability models assume traceable causation. If a defective component causes harm, we can identify the manufacturer, test the component, and establish responsibility.

When an AI agent calls three language models, two computer vision systems, and a proprietary risk assessment API from different vendors, with different model versions deployed at different times, and produces a flawed actuarial prediction, who bears responsibility? Of course, the answer would require a forensic-level investigation that may be

technically impossible after the fact, if the specific model versions, training data states, or API responses were not logged with sufficient granularity, or if the models are too complex.

Viewed through the prism of legal liability, perhaps this is where due diligence and the foreseeability of harm become critical. The reasonableness of the initial deployer's diligence of all parties in the supply chain process may become front and center in assessing blame for harm. As AI systems evolve in sophistication, it is difficult to appreciate what disparate federal and state regulatory agencies might consider when determining ultimate liability, but AI oversight and governance will no doubt be considered.¹⁹

Meaningful validation testing, ongoing monitoring for anomalous outputs, contractual provisions requiring vendors to disclose training data sources and model updates, and detailed logs for post-hoc investigation are factors that may determine accountability and ultimate responsibility. The seminal query may be simple: Should you have been able to detect the symptoms of poisoning, for example, through your validation processes, even if you couldn't identify the source? Of course, this shifts the inquiry from "who is responsible for the model poisoning" to "who failed to catch it!"

Handling your Known Unknowns

A shopper purchasing cookies for a childcare center should check for a nut-free label but is not expected to check its accuracy or the manufacturer's regulatory compliance. Similarly, prior pragmatic, proportional, and practical due diligence, but not forensic inspection of AI systems, may be expected. Unlike cookies, the absence of regulations for AI system labels leaves due diligence up to end users.

"Knowns" are initially established, including the system's intended use, exposed tangible and intangible assets, including processes, information, reputation, and regulatory and contractual obligations. This aligns with established supply chain risk management standards, including the NIST 800-161. The goal is a picture of risks for identified assets to determine justifiable levels of reasonable due diligence.

AI-specific due diligence within existing supply chain risk processes can then be undertaken. Unlike typical supply chain risk processes, the complexity of AI requires distributing questions across distinct roles, where our taxonomy may help.²⁰ The taxonomy identifies the distinct roles of creators, developers, hosts, aggregators, integrators, and users. This enables asking the right questions of the right people regarding the data, models, programs, and infrastructure that AI systems depend on.

For example, analysis may reveal that customer data is being sent to an upstream model host that incorporates it into a public model, thereby contravening regulations or contracts. Contraventions may include inappropriate data use, breaches of privacy or data sovereignty regulations, and failure to de-identify data. Data removal and destruction requirements may become problematic. A cake can't be unbaked, and while cutting a slice brings insights, ingredients can't be completely removed.²¹ Similarly, after most models are trained, analysis might reveal sensitive data, but it's impossible to verifiably remove it.

¹⁹ Frances M. Green, Eleanor Chung, and Raymond Sheh; September 22, 2025; "The Dark Side of AI: Assessing Liability When Bots Behave Badly," *New York Law Journal*; accessed January 8, 2026; <https://www.law.com/newyorklawjournal/2025/09/22/the-dark-side-of-ai-assessing-liability-when-bots-behave-badly/>. Dale Hall and Frances M. Green; November 12, 2025; "Winning the Race — America's AI Action Plan," *Get Plugged In — AI Insights*; Society of Actuaries Research Institute; accessed January 8, 2026; <https://getpluggedin.libsyn.com/>.

²⁰ Raymond K. Sheh and Karen Geappen; November 19, 2025; accessed January 8, 2026; "Identifying the Supply Chain of AI for Trustworthiness and Risk Management in Critical Applications," arXiv preprint arXiv:2511.15763; <https://arxiv.org/abs/2511.15763>.

²¹ Karen Geappen, June 22, 2023, "Cakes Can't Be Unbaked: Why You Should Think Twice About AI," Anchoram Consulting (blog), accessed January 8, 2026, <https://anchoramconsulting.com/au/blog/security/cakes-cant-be-unbaked-why-you-should-think-twice-about-ai/>.

Implications

A better understanding of AI supply chains supports better questions. Who created the training data, and what quality controls are in place? What is the objective function of the training? How to evaluate the model output? How to measure the consistency of the model output? Can model versions be traced when problems emerge? What external data sources does the system access at runtime?

Efforts are underway to extend software supply chain standards to address the unique challenges posed by AI systems. As in other critical sectors such as healthcare and food supply, user and regulatory demand will be crucial to ensuring the adoption of appropriate assurances around its supply chain. Concerned parents succeeded in pushing for “nut-free” labels on cookies. What assurances should concerned actuaries push for their AI systems?

The views and opinions expressed in this essay are those of the authors and do not necessarily reflect or represent the any employer, collaborator, or other entity.

Raymond K. Sheh PhD is Associate Research Scientist at Johns Hopkins University <https://raymondsheh.org>

Frances M. Green Of Counsel at Epstein Becker Green <https://www.ebglaw.com/people/frances-m-green>

Karen Geappen MCSSD is Director Cyber GRC, and Director AI Risk at Anchoram



ACTUARIAL INTELLIGENCE BULLETIN

GenAI Case Study: Leveraging Unstructured Claims Text with LLMs for Improved Cost Prediction

SIMON HATZESBERGER CERA, PHD AND IRIS NONNEMAN MSC AAG

In the article *Advanced Applications of Generative AI in Actuarial Science: Case Studies Beyond ChatGPT*, we demonstrate the transformative impact of Generative AI (GenAI) on actuarial science, illustrated by four implemented case studies. Here we describe case study 1, which shows how Large Language Models (LLMs) can improve claim cost prediction accuracy by deriving significant features from unstructured claim descriptions.

Case Study 1

The insurance industry has long faced challenges in fully leveraging unstructured text data that contains valuable insights for claims assessment and risk management. Traditional analytical methods for claim cost prediction primarily focus on structured tabular data, thereby overlooking critical details embedded in texts such as claim descriptions, incident reports, and customer communications.

This case study demonstrates how LLMs can transform unstructured textual data into structured, actionable information. Using a workers’ compensation claims dataset, we examine how LLMs extract key information such as injured body parts and accident causes from claim descriptions, then use this information to construct new features for a gradient boosting model tasked with predicting compensation costs.

Approach and Techniques

We used a dataset of 3,000 workers’ compensation claims from a Kaggle competition on actuarial loss estimation. This fully synthetic dataset includes both structured features (age, gender, marital status, wages) and unstructured text descriptions of claim reports.

We created a baseline model using existing tabular data as features to predict ultimate incurred claim costs via gradient boosting regression. Next, we employed an LLM to analyze claim descriptions and extract structured information.

To be more precise, the LLM was prompted to extract information following a specific schema: number of body parts injured, main body part injured, and cause of injury specified by verb (see Figure 1). The LLM outputs yielded three additional features. Since the LLM completions produced many different categories for injured body parts and causes of injury, we reduced dimensionality by grouping extracted values through a combination of chatbot suggestions and human judgment.

Figure 1
PROMPT FOR THE LLM TO EXTRACT NEW FEATURES

```

1  prompt = ''' Your task is to extract structured information about injuries and cause of injury from
      the given text.
2  Follow this schema strictly:
3  - number_of_body_parts_injured: The total count of injured body parts.
4  - main_body_part_injured: The primary body part affected, described concisely (e.g., "HEAD", "THUMB
      ").
5  - cause_of_injury: Specify by verb.
6      - If verb given: return only the primary action verb that directly caused the injury (e.g. "fall
      " not "fell from box")
7      - If cause is not mentioned, infer from context if possible, otherwise return "unspecified".
8  Ensure accuracy and consistency in the extracted details. Do not add interpretations beyond the
      provided text. '''

```

We then constructed an enhanced gradient boosting regression model incorporating both the original structured features and the LLM-extracted features. Before fitting both models, we applied a log transformation to the target variable to address distribution skewness and employed grid search cross-validation for hyperparameter tuning.

Results

Our LLM demonstrated good capability in extracting meaningful structured information from claim descriptions. Using human feedback to evaluate LLM completions, we mapped the new categorical features in broader categories to reduce the dimension size.

The enhanced model incorporating LLM-extracted features significantly outperforms the baseline model. *The results show an 18.1% reduction in RMSE and an increase in R² value from 0.267 to 0.508. Moreover, the Mean Absolute Error (MAE) improved by 23.9%.* This demonstrates that the inclusion of features extracted from unstructured text substantially improves predictive quality across the considered evaluation metrics.

Furthermore, our feature importance analysis reveals that the most influential predictors of ultimate incurred claim cost originate from both existing data and newly created claims report information. *Weekly-Wages* consistently ranks among the most significant predictors, with higher wages correlating to increased claim costs. *Age* emerges as another critical factor, with older workers typically associated with higher expenses due to extended recovery periods. Notably, several LLM-extracted features demonstrate substantial predictive power, including *main_body_part_injured* (with important values such as *TORSO* and *HAND_FINGERS*), number of *body_parts_injured*, and cause of injury (with *IMPACT* and *LACERATION* identified as important).

Implications for Actuarial Practice

By leveraging LLMs, we extracted and engineered valuable features from unstructured claim descriptions. These features were grouped into broader categories such as anatomical regions and cause types, which not only improves predictive accuracy but also enhances interpretability by offering insights into drivers of high claim costs.

Our findings demonstrate the added value of integrating LLMs into actuarial modeling: feature importance analysis highlights that both traditional and LLM-derived variables significantly influence outcomes. This allows insurers to improve risk assessment, proactively identify high-cost cases, and design early intervention strategies.

However, care must be taken when incorporating LLM outputs, as their completions can exhibit temporal variance and may require validation. Current research explores systematic evaluation frameworks to enhance reliability and reproducibility. The approach demonstrated has broader applications throughout the insurance value chain, including underwriting, fraud detection, and pricing models. While implementation requires careful consideration of data privacy, model transparency and robustness, and system integration, the potential efficiency gains and improved decision-making capabilities make these techniques a compelling investment for actuaries seeking to transform unstructured data into insightful material.

References

Simon Hatzesberger and Iris Nonneman, “Advanced Applications of Generative AI in Actuarial Science: Case Studies Beyond ChatGPT,” arXiv preprint arXiv:2506.18942 (submitted June 22, 2025), <https://doi.org/10.48550/arXiv.2506.18942>, <https://arxiv.org/abs/2506.18942>

Notebook: https://github.com/IAA-AITF/Actuarial-AI-Case-Studies/tree/main/case-studies/2025/claim_cost_prediction_with_LLM-extracted_features

Dataset: <https://www.kaggle.com/competitions/actuarial-loss-estimation>

Simon Hatzesberger CERA, PhD is Manager Actuarial and Insurance Services, Deloitte
Iris Nonneman MSc AAG is Senior Data Scientist at Achmea



**ACTUARIAL
INTELLIGENCE
BULLETIN**

From Prototype to Production: A Practical Validation Playbook for GenAI in Actuarial Work

CARLOS AROCHA, FSA

Generative Artificial Intelligence (GenAI) can accelerate actuarial deliverables (e.g., management commentary, model documentation, assumption memoranda, etc.), but speed without trust creates an operational, regulatory, and reputational risk. This article provides a lightweight validation playbook actuaries can implement in a few weeks, covering test design, acceptance thresholds, documentation, and ongoing monitoring. And with additional marginal effort, they can map the playbook to recognizable governance expectations (e.g., the AI Risk Management

Framework of the US Information Technology Laboratory,²² NAIC’s Principles on AI,²³ or Switzerland’s Financial Market Supervisory Authority AI Governance and Risk Management²⁴)

Use-Case Scoping

The goal of scoping is to separate “language supplements to actuarial work” (often suitable for GenAI) from “actuarial judgement, calculations, and decisions” (often unsuitable as primary automation). These two questions may help clarify scoping:

1. What role will GenAI play?
2. What will be the consequence if GenAI produces wrong/invalid output?

A simplistic four-tier risk rubric to evaluate whether to use GenAI may be constructed as follows:

Step 1—Assign 0, 1, or 2 points to each of the four dimensions below.

| Dimensions | 0 points | 1 point | 2 points |
|-----------------------|----------------------|--|---|
| External impact | internal only | could influence decisions | client work, regulatory filings, financials |
| Numerical criticality | narrative only | includes figures but can be cross-checked | produces official numbers |
| Data sensitivity | public/non-sensitive | confidential but non-personally identifiable information | personally identifiable information, protected health information, proprietary models |
| Explainability | low | moderate | high evidentiary standard |
| Process coupling | standalone output | feeds downstream work | embedded in operational decisioning |

Step 2—Add the points from the above rubric and use the following criterion:

| Tier | Type | Remark |
|----------------|------------|---|
| A (0—3 points) | low risk | GenAI can draft a light review |
| B (4—6 points) | moderate | GenAI assists, but mandatory cross-checks should be carried out |
| C (7—9 points) | high | confidential but non-personally identifiable information |
| D (10+) | prohibited | do not use GenAI |

A practical decision rule your team can apply in minutes

²² National Institute of Standards and Technology, “AI Risk Management Framework,” Information Technology Laboratory, accessed January 8, 2026, <https://www.nist.gov/itl/ai-risk-management-framework>.

²³ National Association of Insurance Commissioners, “National Association of Insurance Commissioners (NAIC) Principles on Artificial Intelligence (AI),” adopted August 14, 2020, PDF, accessed January 8, 2026, <https://content.naic.org/sites/default/files/inline-files/NAIC%20Principles%20on%20AI.pdf>.

²⁴ Swiss Financial Market Supervisory Authority FINMA, “FINMA Guidance 08/2024: Governance and Risk Management When Using Artificial Intelligence,” December 18, 2024, PDF, accessed January 8, 2026, <https://www.finma.ch/en/~media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmittelungen/20241218-finma-aufsichtsmittelung-08-2024.pdf>.

A use case is a strong candidate when:

- The output is primarily text, not the “official number”
- Every numeric claim can be cross-checked against a system-of-record table
- Errors are caught quickly (review is feasible)
- Inputs can be sanitized (no unnecessary sensitive data)
- You can constrain the task (templates, structured outputs, retrieval from approved sources)

Validation Test Actuaries Can Run

1. Numeric consistency checks (e.g., tables are consistent with the narrative, and vice versa)
2. Source-grounding and citation discipline
3. Hallucination and omission test with “known answer” sets
4. Bias and fairness spot checks for customer impacting text
5. Privacy leakage (e.g., if inputs include personally identifiable information, it is likely that outputs will also include such information)
6. Prompt-injection resilience
7. Reproducibility across runs, including temperature ²⁵ controls and deterministic modes

Creating a One-Page “GenAI Model Card for Actuarial Use”

Think of this as “model governance, adapted to GenAI”: you want clear ownership, reproducible behaviour, an auditable trail, and controlled change without turning it into a months-long program.

Include the following elements in one page:

1. Use case & scope (you may use the rubric provided above or create your own)
2. Risk tier and rationale
3. Model configuration (provider/model version, temperature, tools enabled, retrieval sources)
4. Data handling (describe what is allowed and prohibited, masking rules, storage/retention)
5. Known limitations (failure modes, unsupported tasks, numeric constraints)
6. Acceptance criteria (examples of what “good” looks like; required checks)
7. Approvals (owner, reviewer(s), effective date, etc.)

Other governance artifacts include:

- **Prompt card**—including system prompts/guardrails, approved templated, output format requirements, prohibited behaviours (in particular “inventing citations” or generating official figures)
- **Validation & test evidence pack**—including test set description (representative samples, edge cases), results against the validation tests shown above)
- **Audit-ready logs**—including timestamps, model versions, data classification tags, output versions, etc.

²⁵ In a large language model (e.g., ChatGPT), temperature is a setting that controls how random vs. Deterministic the model’s next-word choices are. In general, low temperatures mean conservative outputs (better for summaries), and high temperatures mean more variety and creativity. For actuarial work, it may be recommended to set the temperature to near zero, to reduce variability.

Conclusion

GenAI can create real efficiency in actuarial teams when it is deployed with clear boundaries: it should accelerate drafting, synthesis, and structured transformation, not replace the systems and professional judgment that produce and defend the numbers.

By scoping use cases up front, applying risk-tiered controls, and maintaining a lean set of governance artifacts, actuarial teams can capture speed gains while preserving traceability, privacy, and accountability. The practical objective is simple: make AI-assisted work easier to review, harder to misuse, and straightforward to monitor as models, data, and expectations evolve.

Carlos Arocha, FSA, is Founder and Managing Partner of Arocha & Associates GmbH. He can be reached at ca@ArochaAndAssociates.ch



The Augmented Actuary in the Age of Agentic AI

BRITTANY LEE, FSA, CERA, AND SHIRLY WANG FSA

The emergence of agentic Artificial Intelligence (AI) is poised to augment the role of the actuary, transforming the practice from manual data compilation and siloed analysis into a high-level coordination and strategic decision-making function. At the heart of this transformation is the concept of the AI agent, a sophisticated system that utilizes a Large Language Model (LLM) as its central “brain” for reasoning, planning, and executing complex tasks. This ability to reason and plan, rather than simply follow a rule-based solution, is what makes agentic AI a truly powerful new desk-mate for the human actuary. In the workshop — Actuarial Innovation Stations: An Expo of Tech & Transformation - hosted at the 2025 SOA ImpACT Conference, we demonstrated an agentic setup to seamlessly integrate the analysis of both unstructured (e.g. actuarial guidelines, insurance documents) and structured (e.g. tabular internal and external data) actuarial data — a necessity for actuarial work that has been complex and time-consuming.

Our demo setup comprises three agents:

- A **Text-to-SQL Agent** to access internal tables, securely housed in the actuary’s ecosystem, and translate the natural language question into precise SQL queries to retrieve specific numerical results
- A **Document Q&A Agent** to access a pool of industry papers, research reports, and other unstructured documents through a vectorized index related to the specific topic
- A **Supervisor Agent** to process a user request, orchestrate and coordinate multiple tasks across the two agents above, and synthesize insights from individual tasks into a single, comprehensive answer

We curated the agent setup to perform analysis for a Long-Term Care (LTC) experience study. Specifically, the Text-to-SQL agent was connected to tables containing LTC incidence and termination data, while the Document Q&A agent was configured to reference a curated set of industry papers on LTC experience. The demo was built using an enterprise AI platform, Databricks Agent Bricks, whose no-code capabilities make it straightforward to stand up the

system. Actuaries are encouraged to develop similar agent capabilities using other no-code or low-code SaaS platforms, and, if a code-centric solution is preferred, by coding with widely-used agent frameworks such as LangChain.

During the workshop, we demonstrated how this agentic system answers common actuarial questions in the context of an experience study. Participants saw how the supervisor agent decomposed the analysis into smaller tasks, issued requests to both the Text-to-SQL agent and the Document Q&A agent, incorporated insights from completed tasks to refine subsequent steps, and ultimately synthesized a comprehensive actuarial analysis. For example, to address the question “Is my incidence data credible by age and gender according to industry convention?”, the agent first used the Document Q&A agent to identify credibility criteria from long-term care incidence studies, then used the Text-to-SQL agent to query claim counts in the specific experience data, and finally compared those results with the industry criteria to deliver a full credibility assessment, much as a human actuary would. The session also included engaging discussions on building trust in AI solutions and avoiding common generative AI pitfalls such as hallucinations.

This agentic approach significantly enhances both actuarial efficiency and the depth of analysis, going well beyond simple task automation. By synthesizing structured and unstructured data, the agent system removes the manual effort of cross-referencing internal data with external and internal documents and supports actuaries with more robust insights. Ultimately, the multi-agent framework elevates actuarial work by transforming previously intractable or time-intensive cross-domain analyses into immediate, actionable insights, paving a clear path toward a more strategic, augmented actuary of the future.

Brittany Lee, FSA, CERA, VP, Marketing Actuary - Partnership Development @ SCOR

Shirly Wang FSA, Delivery Solutions Architect @ Databricks



ACTUARIAL INTELLIGENCE BULLETIN

Do you need AN explanation or THE explanation?

RAYMOND SHEH PHD AND KEVIN MISENER

Start As AI systems increasingly influence legally consequential decisions in insurance pricing, claims processing, and risk assessment, their complexity presents a fundamental challenge: the very capability to learn complex patterns makes these systems so powerful also makes them difficult to trust and verify.

We’re accustomed to models we can audit, validate, and defend. When a traditional mathematical model produces unexpected results, we can trace the calculation step by step. But when a neural network denies a claim or flags fraud, explaining why becomes far more difficult. Even seemingly benign applications, such as using a chatbot to summarize emails and reports, raise trust issues. Can we use those summaries to make important decisions when we can’t trace why they might leave out or misrepresent important details? Can we trust the explanations provided by techniques that offer to “open the black box?”

Following failures in any safety-critical system, regulators expect root cause analysis that identifies not just what went wrong, but why, sometimes down to first principles. When used in consequential settings, like an insurance claim denial or risk assessment, courts will expect litigants using technology to make a consequential decision to

authenticate that system by showing it produces an accurate result under Federal Rules of Evidence (FRE) §901(d). Doing so may require the introduction of expert testimony under FRE §702, which requires that the testimony be “... the product of reliable principles and methods.”

This expectation for how and why a mathematical model produced the results that it produced is increasingly being applied to AI systems, yet many deployed models cannot provide such explanations. Statistical validation, while valuable, can be misleading when systems operate in varied real-world conditions that testing cannot fully capture.

Explainable, transparent, verifiable AI is often seen as a solution to this problem. For critical or consequential applications, it is arguably necessary to obtain an explanation that is both faithful to the underlying decision process and actionable by the recipient. We call this “the” explanation. Unfortunately, most of the techniques currently proposed to improve trust in AI systems do not provide such an explanation.

A Framework for Categorizing Explainability

We find it helpful to categorize AI explainability along three continuous, orthogonal dimensions: Source, Depth, and Scope.²⁶ This framework helps practitioners match explanation requirements to specific use cases.

Source: Where Does the Explanation Come From?

The first distinction concerns whether explanations derive from observing the system's behavior or from examining its actual decision-making process.

Post-Hoc Rationalizations observe an AI system’s outputs for various inputs and construct simplified models to explain them. Shapley Additive Explanations (SHAP) are a well-known example of this.²⁷ These methods are analogous to fitting a function to sample points. This is useful within bounds, but potentially misleading outside them, especially if the underlying process doesn’t match the function. For any finite set of observed behaviors, infinitely many explanations could be consistent with them. These represent **an** explanation, not necessarily **the** explanation.

Introspective Explanations come from systems whose internal representations are inherently interpretable by the target audience. Decision trees that expose their IF-THEN logic, physics-based models where learned parameters have real-world meaning, and some forms of mechanistic interpretability tend towards this type of explanation. These provide **the** explanation for a decision, though user actionability depends on the application.

Limiting models to interpretable representations constrains their complexity and potentially their accuracy. For some applications, knowing the true cause of a decision may outweigh marginal performance gains. In some cases, hybrid neuro-symbolic systems may provide a suitable compromise.

Depth: How Far Does the Explanation Go?

Explanations vary in how deeply they probe the decision-making process:

²⁶ Raymond Sheh and Isaac Monteath, “Defining Explainable AI for Requirements Analysis,” *KI - Künstliche Intelligenz* 32 (2018): 261–266, accessed January 8, 2026, <https://doi.org/10.1007/s13218-018-0559-3>.

²⁷ Scott M. Lundberg and Su-In Lee, “A Unified Approach to Interpreting Model Predictions,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS’17)* (Red Hook, NY: Curran Associates Inc., 2017), 4768–4777.

Attribute Identity explanations indicate which inputs were considered, such as a saliency map showing which regions of an image influenced a decision. For many applications, knowing that the system considered relevant factors may be sufficient.

Attribute Value explanations reveal not just which inputs mattered, but what about them triggered the decision. These might take the form of thresholds ("claim >\$50,000") or examples of cases with different values for that attribute ("here are similar historic cases with claims around \$50,000"). These explanations can be complex for large models that consider many attributes simultaneously.

Model explanations address why the learned model considered particular attributes. This typically requires reference to training data. An example might be: "In 95% of historical cases with these characteristics, the claim amount was important."

Scope: What Can the Explanation Tell Us?

Finally, explanations differ in whether they address specific decisions or teach general principles:

Justification explanations defend particular decisions, including hypothetical scenarios. They answer: "Why did the model do this?"

Teaching explanations convey how the model generalizes, helping users understand where it might succeed or fail in novel situations. They answer: "How does the model think?" beyond specific examples.

Implications

An explanation that does not meet an application's requirements can be worse than no explanation, as it can engender misplaced trust and inappropriate risk management, just as fitting an oversimplified model to data can result in incorrect extrapolations.

Therefore, different applications demand different points in this three-dimensional space. A fraud detection system might require only **post-hoc**, **attribute-identity**, **justification** explanations if the goal is merely to help a human reviewer find statistical differences. In contrast, a system that influences coverage decisions might require **introspective**, **model-level**, **teaching** explanations that withstand legal scrutiny.

An example of this arose in a legal opinion about the definition of "landscaping," in a dispute between a landscaper and his insurer, over coverage for a negligently installed trampoline. A judge deciding the case used ChatGPT's outputs to find that trampoline installation could be covered (though the appeal was denied on other grounds). This initial finding has been questioned, with subsequent studies receiving the opposite answer from ChatGPT, and significant variation in outputs based on minor changes to the prompt, like using quotation marks or adding the word "typically."²⁸ To appropriately trust this decision, any explanation would have had to be **introspective** and **teaching** to reveal such variations.

As actuaries increasingly work with AI systems, understanding these distinctions enables more informed decisions about which techniques are appropriate for specific applications, and clearer communication with stakeholders about what explanations can and cannot deliver.

²⁸ James Grimmelmann, Benjamin Sobel, and David Stein, "Generative Misinterpretation," SSRN (posted June 18, 2025; last revised August 7, 2025), accessed January 8, 2026, <https://doi.org/10.2139/ssrn.5309575>.

The views and opinions expressed in this essay are those of the authors and do not necessarily reflect or represent the any employer, collaborator, or other entity.

*Raymond K. Sheh PhD is Associate Research Scientist at Johns Hopkins University <https://raymondsheh.org>
Kevin Misener JD Technology Attorney*

**ACTUARIAL
INTELLIGENCE
BULLETIN**

How Actuaries Are Engaging With AI: Insights from the SOA's Summer 2025 Survey

JOE ALAIMO, ASA, ACIA

Artificial Intelligence continues to reshape industries across the globe, and actuarial practice is no exception. Yet exactly how actuaries are incorporating AI into their work, and how they feel about its role, has often been a matter of speculation. To establish a factual baseline, the Society of Actuaries Research Institute conducted its inaugural AI Survey in the summer of 2025, drawing 518 responses from across practice areas, geographies, and experience levels.

The goal is ambitious but clear: track AI adoption among actuaries over many years, identify challenges and opportunities, and guide the profession's strategic development in this rapidly evolving area.

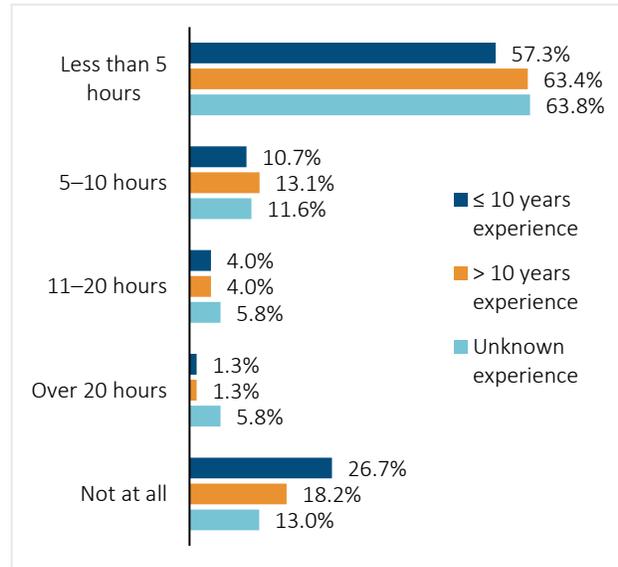
The results offer a nuanced, data-rich picture of where actuaries stand today and, in some cases, the results challenge common assumptions.

The First Surprise: AI Engagement Is Lower Than Expected

One of the first questions the survey asked was simple: How much time do you spend learning about or using AI or ML each week? Despite widespread discussion of AI, most actuaries are not spending significant time with it. More than half of respondents across all experience levels reported less than five hours. Among early-career actuaries with ten or fewer years of experience, 57% fall into this category, and among experienced actuaries with more than ten years of experience, the percentage rises to 63%.

Even more striking, 27% of early-career actuaries reported no engagement at all, compared with 18% of seasoned actuaries. This result runs counter to the common perception that younger professionals are significantly more active with AI tools. In other words, the data shows that senior actuaries actually lead in adoption, although only slightly.

Figure 1
TIME SPENT WEEKLY ON AI/ML TOOLS



What Actuaries Are Using AI For

Among respondents who do use AI, certain applications stand out clearly:

- Learning and brainstorming ideas (50% of early-career respondents and 60% of experienced respondents)
- Writing or interpreting documents (48% and 54%)
- Chatbots (50% and 40%)

These uses reflect the current state of generative AI itself: it performs well at summarizing, structuring, generating text, and supporting exploratory thinking.

More technical applications, while emerging, are less prevalent. Code generation is used by 48% of early-career actuaries, but only 29% of those with more than ten years of experience. This suggests a generational difference in comfort with coding-related tools, although it does not translate into higher overall AI adoption among early-career respondents.

AI Helps, but Mostly by Saving Time

Across experience levels, the top reported benefit is clear: time savings.

- Early-career: 86%
- Experienced: 80%

The next most common benefits include expanded work product and improved decision-making. Experienced actuaries report these slightly more often, which suggests they may be using AI in broader or more strategic contexts.

Accuracy improvements remain the least cited benefit, with only 13 to 19% selecting it. This aligns with themes in the open-ended comments, where many actuaries expressed caution regarding reliability, hallucinations, and the level of validation required before accepting AI-produced results.

Barriers: Compliance, Skills, and ROI

When it comes to organizational barriers to broader AI adoption, three themes dominate:

- Regulatory and compliance risks (approximately 52 to 58%)
- Skill gaps (38% among early-career respondents and 51% among experienced respondents)
- Unclear return on investment (47% early-career and 32% experienced)

Only a small minority indicate that there are no barriers. These results make sense for a profession grounded in risk management. Actuaries recognize the potential of AI, although they also understand the significant responsibility associated with its use.

Open-ended responses highlighted concerns about data confidentiality, accuracy, security, and the need for clear validation frameworks. In many cases, respondents indicated that corporate restrictions prevent experimentation.

Why Non-Users Do Not Use AI

Among the 96 respondents who spend no time with AI, the primary reasons include:

- Perceived lack of relevance (65% early-career; 43% experienced)
- Data quality and privacy concerns (65% early-career; 41% experienced)

Open responses revealed deeper ideological objections as well, including environmental concerns, ethical concerns, skepticism about accuracy, and concerns that AI may not add meaningful value.

Interest in Learning: Another Counterintuitive Finding

Here the data becomes even more surprising. Among the non-users:

- 55% of early-career actuaries reported that they were not interested in learning more about AI applications in actuarial work.
- Only 32% of experienced actuaries reported no interest.

This finding contradicts the expectation that early-career actuaries are more eager to explore new tools.

Limited Organizational Support

Most respondents indicated that their organizations provide:

- Encouragement for self-directed learning (61 to 69%)
- Access to industry resources (29 to 39%)

However, fewer than 30% report any formal training programs. Roughly one-third receive no formal support at all.

This suggests an opportunity for organizations to provide more structured learning pathways as AI becomes more integrated into actuarial practice.

AI's Role Today: Divided Opinions

Open-ended responses regarding AI's role in actuarial practice revealed four major groups:

- 21% strongly oppose AI use in actuarial contexts
- 24% support limited or cautious use
- 25% see meaningful efficiency gains
- 9% describe AI as transformational

This distribution reflects a profession that is weighing enthusiasm against caution, which is fully aligned with actuarial training and mindset.

A Baseline for the Future

This first wave included 518 respondents, with most respondents being highly experienced professionals. More than half have more than 20 years of experience, and nearly one-third have between 11 and 20 years. This demographic profile means that the current snapshot reflects the perspectives of senior actuaries, although future waves may show shifts as newer actuaries progress. The survey will be repeated once or twice annually, creating a valuable longitudinal view of how AI becomes integrated into actuarial work.

The actuarial profession is uniquely equipped to navigate AI's risks and opportunities. The survey findings show a community that is interested, analytical, and deliberate. AI's role will grow, although it will do so under the watchful guidance of professionals trained to manage risk responsibly.

This survey is more than a snapshot. It marks the beginning of the profession's long-term story with artificial intelligence.

Joe Alaimo ASA, ACIA is CEO of Procomp Consulting Inc.



ACTUARIAL INTELLIGENCE BULLETIN

Editors Award: Best Article of 2025

We are pleased to announce that "**The Rise of LLMs in Risk Management**" by **Syed Raza** has been selected as the **Best Article of 2025** by the editors of the *Actuarial Intelligence Bulletin*. This insightful piece explores how Large Language Models are reshaping the industry by enhancing data analysis, pattern recognition, and predictive modeling. From transforming credit risk assessment and regulatory compliance to a groundbreaking case study on a hybrid Generative AI and Reinforcement Learning framework for reinsurance optimization, Raza's work provides a comprehensive roadmap for the future of the field. While addressing critical hurdles such as model risk, data privacy, and the "black box" nature of transparency, the article illuminates how LLMs are poised to revolutionize real-time risk monitoring and core actuarial tasks.

This article appeared in the [May 2025 issue](#) of the *Actuarial Intelligence Bulletin*.

Editorial Committee

Jon Forster, ASA

Dave Ingram, FSA

Ronald Poon Affat, FSA

Jing Kai Ong, ASA

Frank Quan, PhD

Darryl Wagner, FSA

Thank you for reading the January 2026 SOA Research Institute AI Bulletin. We hope you found these insights valuable. Stay tuned for future editions as we continue to explore the evolving landscape of AI and its impact on the actuarial profession. We encourage you to engage with the SOA Research Institute and share your own experiences and perspectives on AI. For questions, comments, and article submissions, contact rpoonaffat@soa.org.

About the Society of Actuaries Research Institute

Serving as the research arm of the Society of Actuaries (SOA), the SOA Research Institute provides objective, data-driven research bringing together tried and true practices and future-focused approaches to address societal challenges and your business needs. The Institute provides trusted knowledge, extensive experience and new technologies to help effectively identify, predict and manage risks.

Representing the thousands of actuaries who help conduct critical research, the SOA Research Institute provides clarity and solutions on risks and societal challenges. The Institute connects actuaries, academics, employers, the insurance industry, regulators, research partners, foundations and research institutions, sponsors, and non-governmental organizations, building an effective network which provides support, knowledge, and expertise regarding the management of risk to benefit the industry and the public.

Managed by experienced actuaries and research experts from a broad range of industries, the SOA Research Institute creates, funds, develops, and distributes research to elevate actuaries as leaders in measuring and managing risk. These efforts include studies, essay collections, webcasts, research papers, survey reports, and original research on topics impacting society.

Harnessing its peer-reviewed research, leading-edge technologies, new data tools and innovative practices, the Institute seeks to understand the underlying causes of risk and the possible outcomes. The Institute develops objective research spanning a variety of topics with its [strategic research programs](#): aging and retirement; actuarial innovation and technology; mortality and longevity; diversity, equity and inclusion; health care cost trends; and catastrophe and climate risk. The Institute has a large volume of [topical research available](#), including an expanding collection of international and market-specific research, experience studies, models and timely research.

Society of Actuaries Research Institute
8770 W Bryn Mawr, Suite 1000
Chicago, IL 60631
www.SOA.org