

ACTUARIAL INTELLIGENCE BULLETIN



September 2025

TABLE OF CONTENTS

Discover the Power of AI at the 2025 SOA ImpACT Annual Conference	2
Actuarial AI Competence Ladder.....	4
SHAP: A Principled Framework for Explainable Machine Learning.....	7
LLMs: You Still Need to Know What You are Doing	9
Will AI Replace Actuarial Models?.....	10
IAA Database of Case Studies.....	12
Automating Reserving with AI	14
The NAIC Artificial Intelligence—Machine Learning Health Insurance Survey	16
Three Lenses for AI in Actuarial Work: Tool. Thought Partner. Disruptor.	18
A Future for Humans.....	21
Actuarial AI Links	23
Editorial Committee	23
About the Society of Actuaries Research Institute	24

Welcome to the September 2025 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 © 2025 by the Society of Actuaries Research Institute. All rights reserved.

Discover the Power of AI at the 2025 SOA ImpACT Annual Conference

JON FORSTER, ASA

November 2–5, 2025 | Seattle, WA

The 2025 SOA ImpACT Annual Conference is your opportunity to connect with fellow actuaries and explore cutting-edge ideas that will shape the future of the profession. This year's meeting in Seattle offers something for everyone—from cutting edge sessions to engaging workshops—and Artificial Intelligence (AI) takes center stage in several high-impact presentations.

Whether you're new to AI or looking to deepen your knowledge, here's a roundup of sessions that showcase how actuaries are applying AI tools, questioning their implications, and preparing for the profession's next chapter.

Session 3I: Building Trust in AI-Generated Data for Actuarial Models and Other AI Applications

Presenter(s): Andrew William Sykes, ASA, FFA

As AI capabilities grow, so does the responsibility to use them ethically and effectively. This session explores the role of synthetic data and how actuaries can ensure trustworthy model outputs in regulated environments. Presenters Matt Nash and Andrew Sykes will highlight transparency, bias mitigation, and developments in the “second generative AI wave” to demonstrate how AI can be applied with integrity in actuarial work.

Key Highlights:

- Assessing AI trustworthiness in actuarial modeling
- Using synthetic data to address data privacy and scarcity
- Exploring ethical implications of generative AI

Session 3K: Great Debate on the Actuarial Use of AI

Presenter(s): Sherry S. Chan, FSA, EA, FCA, MAAA; Ian G. Duncan, FSA, FCA, FCIA, FIA, MAAA; David Schraub, FSA, CERA, MAAA, AQ

This interactive debate tackles the provocative resolution: *“Dependence on AI erodes essential human oversight and professional judgment.”* Two panelists argue for, two against—and attendees are encouraged to consider the tradeoffs between automation and professional discernment. It's a great opportunity to sharpen your thinking, hear diverse views, and engage with peers on this critical issue.

Why Attend:

- Understand the human vs. AI balance in actuarial decision-making
- Develop frameworks for assessing AI integration
- Experience a unique, thought-provoking format

Join the 2025 SOA ImpACT Annual Conference, Nov 2–5 in Seattle, to see how AI is changing actuarial work. Sessions cover ethics and trust in AI, hands-on use of models like TabPFNv2, and practical workshops that take you from data prep to working prototypes. Whether you're new to AI or pushing advanced models, you'll gain usable skills, governance strategies, and a network of peers ready to lead.

Session 4H: Applying Generative AI to Actuarial Datasets

Presenter(s): Brian McGuire, FSA, MAAA

If you're ready to roll up your sleeves and explore real AI tools, don't miss this hands-on look at TabPFNv2—a model developed specifically for tabular data. Learn how it compares with other machine learning models when predicting healthcare costs or mortality and walk away with the know-how to begin applying it to your own datasets.

Takeaways:

- Overview of TabPFNv2 and its actuarial relevance
- Practical AI applications for real-world problems
- Insights from recent research and Kaggle success stories

Session 4E: The AI Enhanced Actuary: How We Are Creating the Most Elite Actuarial Workforce of the Future

Presenter(s): Jon Forster, ASA, MAAA; David Ingram, FSA, MAAA; Joseph Long, ASA, MAAA; Ronald L. Poon-Affat, FSA, FIA, MAAA

This session dives into how the SOA is empowering actuaries with the knowledge and tools to thrive in an AI-powered environment. From data analysis and machine learning to ethical use cases and bias mitigation, attendees will gain a comprehensive view of how to stay future-ready. You'll also get a closer look at SOA's AI-focused resources, including PD Edge+, AI Insights for Actuaries Symposium, the AI Insights podcast, AI bulletin and more.

What You'll Learn:

- Key AI use cases and trends in the actuarial field
- Responsible AI practices and governance frameworks
- Strategies to future-proof your actuarial skillset

We're also offering two dynamic, hands-on workshops where you'll gain practical skills and knowledge. Introduced at the ImpACT Conference in 2024, these workshops were a hit, and we're thrilled to bring them back for 2025.

Workshop 2A: Actuarial Innovation Stations: An Expo of Tech & Transformation

Presenter(s): Jon Forster, ASA, MAAA; Sian Walker, FSA, FIA; Henry Chen FSA, FCIA, MAAA; Igor Nikitin; FSA, MAAA; Dave Ingram, FSA, CERA; Judah Rabinowitz, FSA, MAAA; Ian Blumenfeld; Josh Lauer, FSA, Yukki Yeung, FSA, MAAA

Gain hands-on experience with a variety of AI applications in this interactive, small-group workshop. You'll move through multiple learning stations, each designed to provide real-world skills you can apply directly in your day-to-day work as an actuary. Guided by expert presenters, you'll collaborate, problem-solve, and explore practical ways to integrate AI into your professional practice. Don't miss your chance to step into the role of the actuary of the future.

Workshop 10: AI and the Actuary – A Hands-On Workshop

Presenter(s): Sian Walker, FSA, FIA; Ronald Poon Affat, FSA; Jing Kai Ong, FSA, MAAA; Bruce Rosner, FSA, MAAA; Frank Siu, FSA, MAAA; Jon Forster, ASA, MAAA.

This limited-attendance workshop brings AI to life through an immersive, collaborative problem-solving experience. Attendees will work through the entire AI workflow—from data prep to evaluation—guided by expert facilitators. Expect a fast-paced, engaging format that equips you with both a working prototype and real-world strategies for implementation.

Workshop Perks:

- Real-time application of the full AI process
- Peer collaboration and team learning
- Immediate takeaways to bring back to your organization

Don't Miss Out

With the rapid pace of AI innovation, these sessions offer timely, actionable insights for actuaries at any career stage. Join us in Seattle to dive deep into AI's current and future role in actuarial work—and return inspired, equipped, and ready to lead.

Jon Forster, ASA MAAA, Professional Development Learning & Content Design Director at Society of Actuaries



Actuarial AI Competence Ladder

BY SHANE LEIB, IGOR NIKITIN, JON FORSTER, MICHAEL LEVINE, DAVE INGRAM

The following is the work of a small ad hoc group (the authors) and is not an official statement of the Society of Actuaries or the Society of Actuaries Research Institute. We are sharing this as a Discussion Draft. We hope that others react and continue refining what we are starting here.

This Competency Ladder is intended to help actuaries assess their current AI proficiency, identify growth areas, and support hiring managers in evaluating AI-related skills.

- Identify where they stand and what could be their next steps in their education process around AI.
- To help hiring managers to see what AI knowledge and skills that they need from their new hire.

Basic Level Competency

What They Can Do

- Effectively use Large Language Models (LLMs) for personal productivity tasks such as drafting communications, summarizing complex documents, brainstorming ideas, and conducting preliminary research, leveraging basic prompting techniques.
- **Basic Prompting:** Familiarity with how to formulate clear, concise, and effective prompts to guide LLMs towards desired outputs and to refine prompts iteratively.
- **Application of Actuarial Standards:** The ability to apply ASOPs, principles, and professional judgment to all work involving AI tools and outputs.
- *Primarily uses free account to access LLMs.*

Ready to audit your AI skills? Use this simple ladder to see where you stand. Are you Basic, Intermediate, Advanced, or Expert? Where would you like to be?

What They Know:

- **AI Tool Utility:** Understanding what common AI tools, particularly Large Language Models (LLMs), can generally do (e.g., generate text, generate images, convert text to audio, summarize, answer questions, coding assistance) and their typical applications in a professional context.

- **Probabilistic Nature of AI Outputs:** Comprehension that AI models produce probabilistic, not deterministic, answers, the nature of training process and data used for training and the implications of this for the reliability and accuracy of AI-generated content.
- **Responsible AI Use and Regulations:** General awareness of organizational policies, ethical considerations, data privacy principles, and regulatory considerations concerning the responsible use of AI tools, including what type of information should not be shared with public-facing AI.

Intermediate Level Competency

What They Can Do

- Utilize specialized actuarial AI tools and platforms (fraud detection, underwriting automation, or risk modeling) within their job functions.
- Choose and collaborate effectively with general-purpose AI tools, such as NotebookLM, for enhanced research, document analysis, knowledge organization, and problem-solving.
- Serve as a critical reviewer of AI-generated outputs, ensuring actuarial soundness and professional integrity., critically reviewing and validating AI-generated insights and outputs for actuarial soundness and alignment with professional judgment.
- Communicate with colleagues and stakeholders about the capabilities and limitations of AI-driven insights, preventing over-reliance on AI outputs.
- Contribute to the identification of business requirements for new AI applications in actuarial workflows.
- *Will have one or several paid accounts to access LLMs.*

What They Know:

- **AI Model Types and Insurance Applications:** Understanding the core concepts of common AI model types used in actuarial work and their specific applications in insurance like fraud detection, underwriting, and risk modeling.
- **AI Model Strengths and Limitations:** Knowledge of the specific strengths, weaknesses, and potential biases of various AI models, including scenarios where AI models might be unreliable or require significant human oversight.
- **Advanced Prompting:** Prompt types, categories and techniques (Multi shot, chained prompts, etc.). Setting up AI Assistants. Prompting for complex research and data analysis tasks.
- **Data for Actuarial AI Applications:** Types of data required for different actuarial AI applications, including data governance, quality, privacy and the importance of data preparation.
- **Application of Actuarial Standards:** Knowledge of the relevant actuarial standards of practice, principles, and professional judgment to all work involving AI tools and outputs.
- **Responsible AI Governance and Emerging Regulations:** Understanding of evolving organizational AI governance frameworks, deeper ethical considerations related to AI use (e.g., fairness, bias, transparency), and a detailed awareness of regulations of AI in actuarial practice. Able to contribute to the compliance process.

Advanced Level Competency

What They Can Do

- Design and conceptualize AI solutions for complex actuarial problems, translating business requirements into technical specifications for model development.
- Collaborate with data science teams to design AI applications and provide actuarial insights for model development. Can source expertise needed for a given problem.
- Implement bespoke AI-driven actuarial models using programming languages (e.g., Python) with minimal assistance.

- Conduct advanced validation of AI models, including fairness, bias, robustness testing and cost analysis, to ensure reliability and compliance.
- Build AI-driven automation tools for actuarial workflows (reserving, stress testing, or capital modeling).
- Capable of leading or contributing to AI compliance functions, including documentation, audit readiness, and regulatory reporting.
- Will sometimes use local open weight models and API keys to access LLMs.

What They Know:

- **AI Model Architectures and Applications:** Knowledge of AI architectures and model types used in actuarial practice, including neural networks, XGBoost, SVMs, time-series models, and transformer models for applications like forecasting or text analytics.
- **Advanced Programming Proficiency:** Expertise in programming languages (e.g., Python) and libraries for implementing, developing, and deploying AI-driven actuarial models and automation tools.
- **Advanced AI Model Validation and Explainability:** Understanding of advanced validation techniques (e.g., fairness, robustness, bias testing) and explainability methods required for complex AI actuarial models.
- **Application of Actuarial Standards:** Applying actuarial standards of practice, principles, and professional judgment throughout the AI model development, implementation, and validation lifecycle.
- **Comprehensive AI Governance, Ethics, and Regulatory Frameworks:** Understanding of enterprise-level AI governance frameworks, advanced ethical considerations (e.g., responsible AI development), and regulatory requirements impacting AI solutions in insurance.
- **Educational & Training Resources for their team:** Can identify resources for training and education to meet competency goals.

Expert Level Competency

We think that there are and will be actuaries who are experts. Actuarial AI specialists. Experts will be researching and developing new tech and pushing boundaries of how AI is used—tech and models. New types of models, training. Bleeding edge stuff.

Management and Executive Competencies

Management and Executives responsible for AI decisions should focus on:

- Strategic alignment of AI initiatives with business goals.
- Oversight of AI governance and ethical frameworks.
- Evaluation of AI investment proposals and ROI.

Conclusion

Where do you find yourself on this ladder? When we each look at this list, we notice that we all have holes here and there that we want to fill in to fully achieve our target for our own competency. We see AI as an important and increasing part of actuarial work.

We invite feedback and collaboration to refine this ladder further.

Shane Leib, FSA, MAAA, Actuarial Science Director at University of Notre Dame

Igor Nikitin, ASA, MAAA, Co-founder & CEO, Nice Technologies LLC.

Jon Forster, ASA MAAA, Professional Development Learning & Content Design Director at Society of Actuaries

Michael Levine, ASA, Senior Manager at Express Scripts

Dave Ingram, FSA, CERA is a Blogger and Podcaster on Crossing Thin Ice.

ACTUARIAL INTELLIGENCE BULLETIN

SHAP: A Principled Framework for Explainable Machine Learning

MICHAEL MAYER, DANIEL MEIER, MARIO V. WÜTHRICH

Introduction and Motivation

Modern predictive models, particularly tree ensembles and deep neural networks, often achieve superior predictive accuracy compared to classical generalized linear models. However, their inherent opacity creates significant barriers to trust, regulatory compliance, and model debugging. SHAP (SHapley Additive exPlanations) addresses this critical gap by providing a unified, axiomatically grounded framework that decomposes any model's predictions into additive contributions from individual covariate components. This letter follows our tutorial Mayer et al. [1], and for an introduction to machine learning tools in actuarial science, see Wüthrich et al. [2].

From Shapley Fairness to Model Explanation

The fundamental property of SHAP is that it assigns each covariate component x_j of an input vector $x = (x_1, \dots, x_q)$ a "credit" ϕ_j for a given prediction $\mu(x)$, ensuring that the sum of all credits plus a baseline ϕ_0 equals the model's output exactly: $\mu(x) = \phi_0 + \sum_j \phi_j$. This additive decomposition is mathematically rigorous and theoretically justified through game theory principles. SHAP's theoretical foundation rests on the four Shapley [3] value axioms from cooperative game theory. Translated to machine learning, they are:

1. **Efficiency:** All attributions sum to the total prediction
2. **Symmetry:** Covariate components with identical marginal impacts receive equal attributions
3. **Dummy:** A covariate component that never changes the prediction in any coalition receives zero attribution
4. **Linearity:** Attributions for the sum of two models equal the sum of their individual attributions

Imagine needing to tell a regulator exactly why a policy was declined or cancelled. SHAP gives you that breakdown, step by step.

These axioms ensure that SHAP provides a unique solution for a fair allocation of prediction contributions among covariates. In the machine learning context, each "player" becomes a covariate component, and the "game payoff" is the model's prediction.

A critical choice in SHAP decompositions is the value function to assess the contribution of each coalition to the game payoff. In machine learning application one typically either chooses the interventional (marginal) contribution (marginalized over the population distribution) or the conditional contribution. The latter more realistically accounts for dependence within the covariates, but its estimation is less practical.

Computational Strategies

A naive Shapley computation requires summing over all 2^p subsets, which becomes computationally intractable for large numbers of covariates ($p > 20$). Several strategies address this complexity:

- Monte Carlo Sampling SHAP: Draws random permutations of the p covariates and adds these sequentially to study the marginal prediction changes
- Kernel SHAP: Approximates SHAP values by fitting a local linear surrogate model by minimizing a weighted least-squares problem

- TreeSHAP: Provides an exact dynamic-programming approach for tree ensembles

Empirical benchmarks demonstrate that TreeSHAP can compute attributions for thousands of instances in seconds, while Monte Carlo or Kernel SHAP require minutes to hours for the same workload.

Local and Global Interpretability

SHAP seamlessly bridges individual prediction explanations and global covariate insights:

- Local Explanations: For individual inputs x , SHAP provides additive decompositions that describe how each covariate contributes to the prediction relative to the baseline ϕ_0 . Waterfall plots visualize these attributions, enabling case-by-case root-cause analysis crucial for actuarial applications.
- Global Summaries: By aggregating attributions across many instances i , SHAP enables several powerful visualizations:

SHAP Importance Measures: Average $|\phi_{i,j}|$ across instances i to rank covariate importance

Beeswarm Plots: Display distributions of $\phi_{i,j}$ colored by covariate values, rough description of feature effects

Dependence Plots: Plot $\phi_{i,j}$ against $x_{i,j}$ colored by interacting covariates, detailed description of effects

A key advantage of SHAP is that feature importance rankings remain consistent and don't flip erratically when underlying models change slightly, unlike heuristic measures.

Practical Applications

SHAP serves multiple critical roles in real-world actuarial deployments:

- Model Debugging: Actuaries can identify covariates carrying undue influence, detect data leakage, or uncover hidden biases against protected groups.
- Monitoring and Compliance: SHAP can be integrated into monitoring dashboards to track concept drift by observing shifts in attribution distributions over time. It provides transparent, auditable explanations meeting regulatory requirements.
- Feature Engineering: Dependence plots can reveal sharp non-linear effects or threshold behaviors, suggesting spline transformations or binning strategies. Pairwise interaction analysis can identify interactions worth encoding explicitly in simpler models.
- Stakeholder Communication: Management and customers gain confidence when they see intuitive breakdowns of individual predictions, building trust in statistical models.

Limitations and Best Practices

Despite its strengths, SHAP requires careful implementation:

- Background Dataset Choice: The selection of background data for marginalization influences both the baseline ϕ_0 and all attributions ϕ_j , requiring representative sampling.
- Computational Considerations: Sampling and kernel approaches can be computationally intensive for high-dimensional data, often necessitating pre-selection of representative subsets.
- Conditional SHAP Challenges: While more faithful to dependencies, conditional SHAP demands reliable estimates of joint covariate distributions, which can be challenging in sparse datasets.
- Causal Interpretation: SHAP explains model behavior rather than causation.

Conclusion

SHAP represents a principled, flexible, and widely supported explanation framework that unifies multiple perspectives on covariate importance and local effects. By anchoring explanations in Shapley values, it guarantees fair and consistent attributions that drive insights across virtually any predictive model. From debugging and

regulatory reporting to feature engineering and stakeholder communication, SHAP provides the mathematical rigor and practical utility needed for trustworthy machine learning in high-stakes domains like actuarial science.

The framework's ability to bridge individual prediction explanations with global model insights, combined with its computational efficiency through methods like TreeSHAP, makes it an essential tool for actuaries working with modern machine learning models. As the field of explainable AI continues to evolve, SHAP's axiomatic foundation ensures it remains a reliable and theoretically sound approach for model interpretation and explanation.

References:

[1] Michael Mayer, Daniel Meier, and Mario V. Wüthrich. SHAP for actuaries: explain any model. *SSRN Electronic Journal*, 2023. Manuscript ID 4389797.

[2] Mario V. Wüthrich, Ronald Richman, Benjamin Avanzi, Mathias Lindolm, Marco Maggi, Michael Mayer, Jürg Schellendorfer, Salvatore Scognamiglio. AI tools for actuaries. *SSRN Electronic Journal*, 2025. Manuscript ID 5162304. <https://aitools4actuaries.com/>

[3] Lloyd S. Shapley. A value for n -person games. In Harold William Kuhn and Albert William Tucker, editors, *Contributions to the Theory of Games (AM-28), Volume II*, pages 307–318. Princeton University Press, December 1953.

This article is an overview of a more complete treatment of this topic: SHAP for Actuaries: Explain any Model (March 15, 2023). Available at SSRN: <https://ssrn.com/abstract=4389797>

Michael Mayer, Actuarial Department, La Mobilière, Bern

Daniel Meier, Swiss Re Institute, Zurich

Mario V. Wüthrich, Department of Mathematics, ETH Zurich

ACTUARIAL INTELLIGENCE BULLETIN

LLMs: You Still Need to Know What You are Doing

STEPHEN MILDENHALL

I often see complaints that Large Language Models (LLMs) aren't useful. That hasn't been my experience at all. I find them fantastically helpful. But the way you interact with them can lead to radically different outcomes.

In my last post, I argued that domain expertise separates professions that are more likely to be amplified by technology (like actuary's post-pcs) from those that risk being replaced (like typists). Here's a concrete example of how domain expertise can help you get what you want out of LLMs.

I recently had to solve the following problem. Take a list of names like:

Smith; smith, m.; smith, michael; smith, michael j.; smith, jane

And map the shorter forms to the most complete version, so smith, m. Becomes smith, michael j. Note that we can only do this where there's a unique, unambiguous match. Smith alone could refer to more than one person, so it stays unresolved.

I asked an LLM to code it up. The first few attempts didn't come close.

Rather than trying to hack it further, I stepped back. Writing software often comes down to combining existing, well-established solutions. You rarely need to rewrite foundational code. Libraries like pandas, numpy, and matplotlib exist for a reason.

So, I asked if there is a data structure that could solve my problem? "Yes," the model replied. "a trie."

I hadn't heard of tries. They're commonly used in auto-completers, I learned. The LLM taught me how tries work, pointed to the pygtrie package, and coded a working prototype. Then I asked the critical follow-up: can we use tries to solve the name-matching problem? Unsurprisingly, the response was an emphatic yes. The next iteration was a clean, efficient, and effective solution.


That's the difference domain intuition makes.

LLMs don't replace expertise. They amplify and enable it, just as we saw in the previous post about typists and actuaries. If you bring structure, insight, and a sense of what kind of solution should exist, these tools can be remarkable collaborators.

But if you're "vibe coding" without a solid grasp of the programming language, the problem domain, or how computing structures work, you're unlikely to fly.

Is your experience the same? How do you make LLMs work best for you?

Stephen J Mildenhall, PhD, ASA, FCAS, MAAA retired blogger at <https://blog.mynl.com/>



Will AI Replace Actuarial Models?

IGOR I. NIKITIN, ASA, MAAA

Will AI replace traditional actuarial modeling? Will we be able to ask a ChatGPT-like tool to run a quarterly valuation or prepare an annual statutory report? In short, I don't think so, at least not for the next decade. AI will certainly make modeling more efficient, but it will not replace traditional modeling in the foreseeable future. Here are the reasons why.

1. AI Models Are Inherently Less Cost-Effective

Traditional models have targeted scopes and are relatively affordable. The current generation of Large Language Models (LLMs) are quite expensive to operate because of high computational overhead cost. Although it may shrink with technological advances, the cost will likely remain non-trivial.

For instance, running a traditional model evaluating 1 million policies might cost about \$10 on a high-end AWS machine. In contrast, we can estimate the cost of running the same model in ChatGPT by calculating the cost of input and output tokens and assuming we get a perfect result. Even under optimistic assumptions, using an LLM

like ChatGPT to evaluate 1 million policies will cost between \$36,000 and \$12 million per run, depending on complexity and the number of intermediate calculations required.

This is significantly more expensive than hiring a pair of actuaries, paying for an actuarial software license, and running the model for \$10 per run.

2. Explaining AI Output Requires a Traditional Model

AI outputs for complex tasks like actuarial modeling are hard to verify due to their opacity and the sheer volume of calculations. How can you be sure the AI made no mistakes? You'd want to see the formulas and the transformations from input to output. This is exactly what a traditional model provides. But if you need a traditional model to verify the AI results, why use the AI model at all, given its high operating costs?

Current LLMs are far more costly at scale, lack the transparency needed to verify complex calculations, can't hold large in-force blocks in context, and struggle with hard math. For now, the practical approach is a hybrid one—use AI to assist modelers, not to replace them.

This problem will only be solved when society and the actuarial profession become comfortable fully trusting AI. We are still very far from this.

3. Current LLMs Have Insufficient Context Awareness

The current generation of LLMs can handle a context of about 128,000 tokens, or roughly 43,000 floating point numbers. Inputs, outputs and intermediate reasoning must all fit within this limit, or the model will produce results without full awareness of the data. In practice, this means that the current generation of LLM's can only handle 1 to 10 policies at a time or require significant additional engineering to limit the data being used for each policy. This effectively replaces the complexity of an actuarial model with the complexity of AI data retrieval infrastructure. That is a rather pointless trade.

To perform a valuation-type run, an LLM would need a much larger context. To handle the entire 1 million inforce block at once, we would need a context that is 15,000 to 20,000 times larger than today's best model. The complexity of LLM processing grows roughly with the square of the context size. Doubling the context quadruples complexity, so increasing it 15,000 times would be about 225 million times more complex, and likely 225 million times more costly.

We need drastically better algorithms and computer power. Something on the order of magnitude of practical quantum computing would help.

4. The Current Generation of LLMs Struggles with Complex Mathematics

The current generation of LLMs is not great at mathematical problems. A 2024 paper by Apple demonstrates and concludes that "the performance of LLMs deteriorates as question complexity increases." There are attempts to solve this issue by pairing LLMs with formal programming languages, such as Google DeepMind's AlphaProof, which recently solved 4 out of 6 International Mathematical Olympiad high school problems. However, it uses a more complex iterative training approach and is much more expensive to operate than ChatGPT, which we used as the basis for our cost analysis here.

Conclusion

To cost-effectively replace traditional actuarial modeling with AI, the AI would need to reliably produce results at a cost lower than both the salaries of actuarial modeling staff and traditional cloud computing expenses. This is extremely unlikely within the next decade, as it requires major improvements in AI algorithms and computing

power. For the foreseeable future, traditional models developed with AI assistance remain the most practical and cost-effective approach to actuarial modeling.

Even if AI eventually replaces today’s actuarial models, the “actuary in the loop” will remain essential. Actuaries must establish and validate models, frame the challenges they address, and interpret the results. They will also be accountable to their principals, the profession, and the public for the integrity of work performed, including portions performed or assisted by AI. This makes it critical to advance both AI-based modeling and actuarial skills in parallel.

Full analysis is available at: www.nicetechnologies.com/will-ai-replace-traditional-actuarial-modeling/

Igor Nikitin, ASA, MAAA, Co-founder & CEO, Nice Technologies LLC.



IAA Database of Case Studies

SIMON HATZESBERGER AND DAVID SCHRAUB

The International Actuarial Association’s (IAA) AI Task Force has developed a rich and growing set of resources for actuaries interested in artificial intelligence. One of the most valuable hubs for these materials is the Task Force’s GitHub account – a central place where actuaries find practical case studies, how-to guides, and tools that help actuaries apply AI in their work.

The IAA AI Task Force

The IAA is the worldwide association of professional actuarial organizations, with a mission to advance scientific knowledge and the skills of actuaries globally. Launched in 2024, the AI Task Force brings together about 100 volunteers from all over the world. In its 2025/2026 phase, the AI Task Force is organized into four workstreams: Engagement and Foundations, Research and Advancement, Case Studies and Tools, and Adoption Framework. This article offers a closer look at the work of the Case Studies and Tools workstream. More information on the AI Task Force – its workstreams, overall deliverables, and general background – is available on the IAA AI Task Force website <https://actuaries.org/council/executive/artificial-intelligence/>.

If you have visited the AIforActuaries platform, you’ve already encountered one of the AI Task Force’s flagship initiatives – a space for learning, discussion, and exchange of AI-related content tailored for actuaries.

To share resources on case studies and tools more broadly, the AI Task Force also maintains a dedicated presence on GitHub, a widely used platform for hosting and collaborating on code, documentation, and project materials. The GitHub account features two primary repositories:

- Actuarial AI Case Studies: a curated collection of real-world AI examples – including machine learning, Generative AI (GenAI), Agentic AI—designed to help actuaries apply these techniques to practical problems.
- AI Tools for Actuaries: a comprehensive hub of tutorials and how-to guides on AI tools, from AI-assisted meeting notes to usage of Large Language Models via API, aimed at enabling actuaries to integrate AI technologies into their workflows.

In the following sections, we take a closer look at each repository, explaining their purposes and showcasing sample content.

Repository Spotlight: “Actuarial AI Case Studies”

This repository compiles a wide range of actuarial case studies involving AI, covering all skill levels from introductory to expert. Each case study is listed with metadata (e.g., date, topic, complexity level, models used, abstract/summary) and includes links to supporting materials such as articles, websites, and code. The collection primarily compiles existing case studies from the actuarial field; however, for topics where no suitable examples were found, the AI Task Force developed original case studies – complete with implemented code (often in notebooks) and explanatory documentation.

The repository’s purpose is twofold:

1. Demonstrate concrete AI applications to actuarial tasks.
2. Enable actuaries to adapt these solutions for their own work by providing code and practical implementation steps.

To illustrate the repository’s value, we highlight two original case studies showing how GenAI streamlines actuarial tasks with complex visual and unstructured data. Both are fully implemented as Jupyter notebooks.

Case Study 1: Car Damage Classification and Localization with Fine-Tuned Vision-Enabled LLMs

This case study explores how multimodal Large Language Models (LLMs), such as OpenAI’s GPT-4o, can classify different types of car damage from images. It compares three models: a traditional Convolutional Neural Network (CNN), an off-the-shelf GPT-4o, and a fine-tuned version of GPT-4o trained on a labeled subset of the car damage images. For the given task, the off-the-shelf GPT-4o performed nearly as well as the CNN straight out of the box – benefiting from pre-training and without requiring specialized CNN expertise – while fine-tuning the LLM further enhanced its accuracy and F1 score, surpassing the CNN’s results.

Notably, GPT-4o can go beyond simple classification by extracting contextual insights, such as the damage location (e.g., “rear bumper”) and additional details like weather conditions or license plates. These capabilities are particularly valuable for automating claims processing and risk assessment.

Case Study 2: GenAI-Driven Market Comparison

This case study applies GenAI to extract and compare key financial and risk data from insurance company annual reports – documents that are long, unstructured, and highly variable. Using a three-stage pipeline based on Retrieval-Augmented Generation (RAG), the model extracts specific aspects such as solvency capital ratios, interest rates by duration, and cyber risk mitigation strategies, demonstrating its ability to handle diverse types of data.

The pipeline first processes and embeds the documents for semantic search, then retrieves relevant text chunks based on the user’s query and finally uses Structured Outputs to ensure the results conform to predefined JSON schemas. This approach enables actuaries to efficiently produce standardized comparisons across companies without manually reviewing entire reports – and can be readily adapted to analyze and compare other documents, such as risk reports or insurance tariff details.

Repository Spotlight: “AI Tools for Actuaries”

This repository offers practical tutorials and how-to guides on AI tools designed to support actuaries in their daily tasks. Content is delivered in formats best suited to each tool – whether short instructional videos, interactive notebooks for code-based applications, or other appropriate media. Examples include AI tools that automatically extract meeting notes and action items from online meetings, applications that assist with programming tasks such

as migrating from legacy systems to modern platforms, and guidance on using LLMs via API, including basic prompting as well as advanced GenAI techniques like Function Calling and RAG.

The repository’s purpose is to raise awareness of relevant AI tools, demonstrate their practical use with best practices, and highlight how these technologies can improve the efficiency and/or quality of actuarial work.

Get Involved: Grow the Actuarial AI Community

We invite you to explore the IAA AI Task Force GitHub repositories and discover practical ideas to adapt to your work. If you know of relevant case studies or tools – whether developed by you, encountered in practice, or worth sharing – please submit them using the repository guidelines or contact us directly. Help spread the word to colleagues so that more actuaries can benefit from and contribute to these resources.

New content is added regularly – stay connected and check back often for the latest updates!

Simon Hatzesberger, Aktur DAV, CERA, PhD is Actuarial and Insurance Services Manager at Deloitte
David Schraub, FSA, MAAA, CERA, ACA, AQ is Founder and CEO of david schraub actuarial consultancy



Automating Reserving with AI

INTERVIEW WITH OLI GROSSMAN

AIB: Oli, AI has started to reshape many traditional fields. How do you see this shift impacting actuarial science?

Oli Grossman: We're seeing AI come in from a lot of angles, be it algorithmic underwriting; image recognition in claims or new modelling techniques in reserving. Actuarial work has always relied on statistical models, but recent advancements in AI have allowed us to begin automating even more complex processes.

AIB: You've been at the forefront of this shift, with the development of a major actuarial application. Can you tell us about it?

Grossman: Of course. We created an AI-backed reserving platform. It’s able to automate the entire reserving workflow: spotting trends in key reserving diagnostics; setting development patterns automatically; and flagging the main sources of uncertainty. From there, you can refine the model’s suggested assumptions and set your final reserves. As reserving actuaries, we know we’ll never be perfectly “correct,” so we designed a platform that uses AI to produce insights you simply couldn’t otherwise get and help you better understand your claims-portfolio.

AIB: Can you walk us through how AI improves the reserving process specifically?

AI is reshaping actuarial reserving by automating time-consuming tasks and surfacing insights humans can miss. Oli Grossman describes an AI-backed platform that detects emerging trends across portfolios, sets assumptions and tail curves automatically, flags the main drivers of reserve uncertainty, and speeds up capital model checks — while keeping a human in the loop for oversight and communication.

Grossman: Reserving can be labor-intensive. For example, analyzing claims data, applying triangle-based methods and searching for emerging trends all take time. With AI, we're able to automate these steps and transparently too. For example, we can:

- Identify emerging trends across all segments and aggregations of a portfolio.
- Flag the main assumptions driving reserve uncertainty.
- Automate assumption setting, e.g., Initial Expected forecasting, age-to-age exclusions, and tail curve fitting.

These allow us more time for value-add judgement and reduce the need for data wrangling & creating diagnostics.

AIB: How is AI being used beyond individual insurers—for example, at the systemic level?

Grossman: That's an area where AI has gained ground quickly. Regulators and oversight bodies are now using machine learning to assess reserve adequacy across insurers. These tools can flag sectors, companies or lines of business that are showing early warning signs of weakening reserves, changing claims behavior or unaccounted inflation.

AIB: That sounds like it could also bring more consistency to regulatory oversight.

Grossman: Yes – that's exactly right. These kinds of reporting tools, where a standardized process is run repeatedly and reproducibly, provide more standardized frameworks for evaluating reserve adequacy or whatever is being assessed. That consistency means actuaries and regulators can compare trends across insurers over time with greater confidence and transparency.

AIB: What about economic capital model validation? How is AI changing the way that's done?

Grossman: Capital models require extensive validation to ensure they are modelling insurers' risk profiles as expected. So, with a similar mindset, we've developed a platform that automates much of this process. It runs hundreds of validation schedules across a wide range of parameter-inputs and can even be extended to incorporate AI techniques for deeper pattern recognition and predictive modelling. This makes validation more efficient and more robust—again freeing actuaries to focus more on interpreting and communicating results instead of putting them together.

AIB: There's a lot of talk about whether AI will replace certain professional roles. What's your take on that, especially in the actuarial profession?

Grossman: I don't think the integration of AI into actuarial work will replace actuaries...yet. What I do think will happen in the next 1–3 years is that we will have smarter tools at our disposal to navigate an increasingly complex risk landscape. The key thing to bear in mind in our profession is that actuaries still need to ensure that their work aligns with business needs, meets regulatory standards and is well communicated to the functions that comprise an insurance company. Most AI-driven tools we use cannot do this reliably and therefore still need to operate with a human-in-the-loop.

AIB: Looking to the future, where do you see AI making its next big impact in actuarial science?

Grossman: I think that the next wave could involve bringing insurance functions together. As data gets captured in more robust and structured ways, actuaries, underwriters and claims teams can all start to plug into the same AI-backed frameworks and begin speaking a common language. One new area we've started developing in light of this is having an ML-driven index of case-reserving strength which quantifies shifts in claims experience - giving all functions a consistent metric to track claims across the business, all feeding back into each other in real time.

AIB: Final question—what’s the key to successful AI adoption in actuarial work?

Grossman: In my opinion, the key is to avoid shifting the goalposts too drastically at the outset. Start by working with what the business already knows but making a few operations smarter—embedding AI into what’s familiar, then pivot gradually. That steady, incremental approach can really help maintain trust in professional work and reduce the barrier to adoption.

Oli Grossman, FIA C.Act is a Senior Consultant at LCP in London

ACTUARIAL INTELLIGENCE BULLETIN

The NAIC Artificial Intelligence—Machine Learning Health Insurance Survey

DOROTHY L. ANDREWS, PH.D., ASA, CSPA, MAAA

The National Association of Insurance Commissioners (NAIC) is committed to understanding how insurers are deploying artificial intelligence (AI) and machine Learning (ML) techniques to provide and service insurance products free of unfair discrimination. Private passenger auto (PPA) insurance was the first insurance line to be surveyed followed by homeowners (HO) insurance, life insurance, and finally health insurance. A major distinction among these lines of insurance was how machine learning was defined and whether generalized linear models (GLMs) are considered machine learning. This debate rages domestically and abroad, at least among actuaries.

Linear regression models have been considered a supervised form of machine learning by computer scientists long before their adoption as a “traditional” ratemaking technique by property and casualty actuaries. The NAIC surveys did not settle the debate as whether GLMs should or should not be considered artificial intelligence, of which machine learning is a subset. For the private passenger auto and homeowners surveys, regulators excluded GLMs from the definition of ML, not because they believed it was not a form of AI, but because GLMs are widely used in PPA and HO insurance and are more of a “glass box” for regulators. Regulators wanted to focus on the more “black box” forms of machine learning techniques insurers are deploying in the PPA and HO surveys.

In the life and health surveys, GLMs were included in the definition of artificial intelligence. If GLMs were excluded from the definition of artificial intelligence in the life survey, there would have been little to study, as life insurers are still nascent in their use of GLMs, with other more advanced techniques yet to be explored. Health insurers, while lagging behind PPA and HO insurers in their use of GLMs, were found to be ahead of life insurers in their use of GLMs. It was surprising that some health insurers do not consider GLMs to be artificial intelligence and advocated to have them excluded from the definition of artificial intelligence used in the survey.

The health survey focused on four product lines:

- 1) Individual Major Medical,

- 2) Group Major Medical Single Employer—Small Employer,
- 3) Group Major Medical Single Employer—Other Employer, and
- 4) Student Health.

These lines were chosen because they fall under state regulatory authority, unlike some plans under the Affordable Care Act. Ninety-three (93) health carriers were selected to participate in the survey, and it was found that 92% of them were using, planning to use, or exploring the use of AI in their operations, which was more than the same percentages for PPA (88%), HO (70%), or life (58%) insurers. It should be noted that considerably more carriers were selected to be surveyed in PPA (193), HO (194), and life (193) insurance.

The health survey consisted of over 500 questions specific to each of the four product lines, plus several more focused on AI testing and governance. The product specific questions surveyed the use of AI in ten (10) insurance operational areas:

- 1) Product Pricing and Plan Design,
- 2) Claims Adjudication,
- 3) Prior Authorization,
- 4) Utilization/Severity/ Quality Management,
- 5) Risk Management,
- 6) Risk Adjustment,
- 7) Fraud Detection,
- 8) Data Processing,
- 9) Sales & Marketing, and
- 10) Strategic Operations.

Insurers were asked to not only indicate whether that AI was being used, but also to indicate the stage of implementation (research, proof of concept, prototype model, implemented in production) and describe the purpose and machine learning technique deployed for each implementation. Companies were also asked to disclose the third-party data and vendors supporting each AI implementation.

The results revealed over 100 different third-party products being employed by health carriers. Health insurers, unlike the PPA, HO, and life insurers, were found to be heavy users of large language models (LLMs) such as ChatGPT 3.5, ChatGPT 4.0, Gemini, Anthropic, and Bart. Machine learning techniques such as Classification and Regression Trees (CART), Cluster Analysis, Convolutional Neural Networks, Decision Trees, Gradient Boosting Machines, Principal Component Analysis, and Logistic Regression, among others, were listed.

A significant number of third-party vendor models were disclosed as well. In fact, insurers indicated that 55% of their models are developed internally with third-party components, and 28% are either developed by a third-party or with the help of a third-party. These are significant findings because currently third-party vendor products are not regulated under statutory insurance provisions, although this is expected to change. Regulators have put the onus of ensuring third party vendors and data do not unfairly discriminate against consumers on insurers. Insurers have expressed a preference to have third party vendors regulated to relieve them of the burden of interrogating third-party model and vendors for unfair discrimination, and some have contractual language in place asking vendors to confirm the models are compliant with insurance laws and regulations. To the delight of regulators, more than eighty-five (85%) of insurers responded that their contracts with third parties do not include any conditions that would limit disclosure or otherwise limit transparency to regulators.

The NAIC's 2025 health AI survey found 92% of surveyed health insurers are using, planning, or exploring AI—more than in auto, homeowners, or lifelines. Health carriers rely heavily on LLMs and varied ML methods; 55% of models are built internally with third-party components and 28% are third-party. The report highlights vendor disclosure, bias-testing approaches (like BIFSG for inferring race), a five-step model-drift monitoring practice, and that nearly 92% of insurers have AI governance aligned with NAIC principles.

Bias testing is dependent upon having protected class attributes. The survey revealed health insurers are employing several methods to infer race. The method disclosed most often was the Bayesian Improved Firstname Surname Geocoding (BIFSG). The best-in-class use of the BIFSG method was discussed in a Three-Phase Waterfall approach, where clinical and administrative data is used as a first step, followed by person-level or family-level imputation processes when clinical data is not available. The BIFSG method is then only employed if the first two approaches fail to identify race. The waterfall approach appears to have a high level of accuracy, and is unique to health insurers, as there is no clinical data equivalent in PPA, HO, or life insurance.

A similar best practice was discussed for managing model drift using a five-step process. The steps include: 1) Monitoring the data and models, 2) Reviewing SHapley Additive explanation (SHAP) values, 3) Conducting bias and fairness evaluations, 4) Continuous testing of model generalizability, and 5) Conducting impact reviews. It is important to continuously check models for drift as changes in human behavior, technology, and the external environment can impact the fit and relevance of a model.

Finally, the survey sought to understand the extent to which insurers were adhering to the NAIC AI Principles of:

- 1) Fairness and Ethicality,
- 2) Accountability,
- 3) Regulatory Compliance,
- 4) Transparency, and
- 5) Security, Safety, and Robustness.

Nearly 92% of the insurers surveyed responded “yes” that they have AI governance principles in place that model the NAIC AI Principles.

Readers are encouraged to visit the NAIC website at the following link to review the executive memo and report on the health survey: <https://content.naic.org/committees/h/big-data-artificial-intelligence-wg>. This article barely scratched the surface of the content contained in the full report. The NAIC AI Principles can also be found at the link above.

Dorothy L. Andrews, Ph.D., ASA, CSPA is Senior Behavioral Data Scientist and Actuary - NAIC.



Three Lenses for AI in Actuarial Work: Tool. Thought Partner. Disruptor.

AREE BLY

AI isn't optional. It's here, and it's reshaping the actuarial landscape. At recent SOA events, conversations ranged from using Generative AI in modeling to grappling with professionalism and privacy, to creating AI Agents to calculate premium rates.

AI is here — and for actuaries it's not just a tool or a threat but three different lenses: a Tool that speeds code, analysis, and writing; a Thought Partner who challenges assumptions and sparks new scenarios; and a Disruptor that changes the data and behaviors our models rest on. This article walks through each lens with practical examples and clear guardrails, so you can adopt AI to gain efficiency and insight while still validating results, spotting upstream changes, and preserving professional rigor.

How we think about and use AI depends on the lens we are looking through. When we view AI as a tool, a thought partner, or a disruptor, we see three distinct ways that it may impact actuarial work.

Lens 1: AI as a Tool

The most immediate and practical lens is using AI to enhance the work that we do today. Like upgrading from a manual screwdriver to a power drill, AI tools can improve both speed and accuracy of certain tasks and make previously difficult work easier.

On the technical side, AI can generate code faster and with fewer syntax errors. It excels at pattern recognition within large datasets, being able to digest and summarize large quantities of data and even identifying subtle correlations or outliers that may be missed in manual analysis.

For actuaries that want to improve their written communication, both technical and otherwise, AI can be a strong editor. It can transform complex technical findings into clear executive summaries, restructure reports for different audiences, and improve readability and conciseness in everyday emails.

In each of these areas, the actuary remains the craftsman and primary analyst. AI may generate code, suggest potential patterns, or revise your writing, but your professional judgment considers whether the code is appropriate, the insights are meaningful, and the communication serves the intended purpose effectively.

Viewing AI through this lens:

- Start with minor routines and tasks to build trust and confidence.
- Peer review the work as you would analysis coming from a competent intern.
- Clearly document what work and content was AI-assisted versus human-produced.

Lens 2: AI as a Thought Partner

Viewing AI as a collaborator with perspectives and access to information beyond your own experience or organizational knowledge, offers an opportunity for greater innovation and creative problem solving.

When new regulations are released, AI can help you think through downstream implications in your organization, actuarial processes and models, and client outcomes. It can help to identify areas in your blind spots, make connections and even offer potential unintended consequences.

It can offer counterarguments and alternative scenarios. For instance, when modeling pandemic-related mortality, AI might suggest considering variables such as vaccination rates, population density, or healthcare system capacity. It can pose questions to you that get you thinking about new angles and implications.

Beyond the technical work, it can offer a space to explore leadership challenges or plan difficult conversations with more empathy and creativity.

This collaborative approach requires deliberate training of AI tools. You aren't looking for positive reinforcement, but seeking out other perspectives, constructive criticism, or a devil's advocate.

Using this lens, you might:

- Have it role play and act like a regulator when reviewing your draft rate submissions and documentation.
- Draft peer review questions around fairness, relevance, and risks.
- Generate scenarios to make you think one level deeper, such as “What if AI triage increases telehealth utilization?”

Lens 3: AI as a Disruptor

While the first two lenses are looking at ways that we, as actuaries, can use AI, we are not the only ones using it. When AI is used by consumers, providers, or even regulators, it changes the data, assumptions, and industry environment.

Perhaps the most actuarial lens that we can take on is to consider how the use of AI is rippling through our work. It is transforming consumer behavior, changing claims processing speed and accuracy, and reshaping business operations and data upstream of our models and analysis.

This means that we need to reexamine our reliance on historical data and adjust the assumptions we use for our projections. We need the actuarial eye more than ever.

In healthcare, AI-driven diagnostic tools are improving early disease detection, AI supported telemedicine platforms are changing care delivery patterns, and predictive algorithms are driving preventive outreach. These innovations alter morbidity patterns, utilization rates, and claims costs.

In life insurance, AI-powered wellness apps may influence policyholder behavior, wearable devices offer real-time monitoring, and consumers use AI to compare policies or understand long-term financial trade-offs.

The key question for actuaries becomes “Where is AI adoption in adjacent industries affecting the fundamental assumptions and data underlying my work?”

This requires regular monitoring of technological trends in your market and an openness to revisit assumptions and models impacted by AI’s ripples.

Through this lens:

- Treat AI as another factor to be understood and considered in your actuarial work.
- Use the Tool and Thought Partner lenses to identify and respond to upstream AI uses.
- Invest time in scenario testing that includes AI-driven disruption.

Maintain Your Actuarial View

What would an actuarial take on AI be without a few limitations?

Effectively integrating AI into actuarial work means balancing our core professional responsibilities with the opportunities that these new tools. This includes:

- Recognizing the limitations and biases inherent in AI tools and output.
- Validating and reviewing AI-generated content before using it to inform analysis or decisions.
- Ensuring analytical rigor is included alongside the efficiency gains.

AI as we know it today is not infallible. It can hallucinate data, perpetuate biases, and produce confident yet incorrect ideas. Our role demands skepticism and clarity, especially when using tools that can do so much, so fast.


Final Thoughts

AI isn't replacing the actuarial mind, but it can extend it. Seen as a Tool, it can lighten the load. As a Thought Partner, it can spark innovation. As a Disruptor, it is a factor to consider in your work.

By using these lenses, we can redirect our actuarial eyes to higher-value activities such as strategic thinking, ethical considerations, and complex problem solving that requires deep domain expertise and creativity.

The future of actuarial work with AI has already begun. You get to decide how you look at it. Using all three lenses together may just become your new actuarial superpower.

Aree Bly is Professional Trainer and Coach at Alignment Ally, LLC



A Future for Humans

DAVE INGRAM, FSA, CERA

Personally, I do not subscribe to the current myth that AI will soon become powerful and competent enough to take over a large fraction of human jobs. But many people do believe that myth or at least fear it. So regardless of the likelihood, it becomes important to prepare ourselves for its implications.

What we need is our own narrative that we can use as a guide that runs along with and counter to that AI myth. A narrative that has us humans as major characters, not just as also rans. Our story will say that we are valuable contributors, and AI is not as stand-alone valuable as its creators claim. The real power of AI will manifest when it is shaped by and combined with human intelligence.

There are four important elements to that narrative.

We need to know our own value proposition. Many of the things that AI can automate may have been part of our jobs, but they are not our real contribution. The fact that you come into the office (or sit down in your home office) and spend an hour answering emails does not make email answering a major part of your value proposition. An actuary's value comes from the results that we can create that cannot be created by others. In general, actuarial value comes from how we can deal with uncertainty, our proven ability to design and manage very long term financial security systems, Our capability to build and evaluate the applicability of models, how we are able to maintain trust of our stakeholders by performing these and other complex financial juggling acts within a framework of rigorous professional standards and our continual focus on what is really important to the long term future of our employers and clients.

We need to know the real value proposition and the real limitations of artificial intelligence. From my own experience with a dozen or more LLMs, I would say that AI can access an extremely wide range of information. It can respond to my statements as if it were a conversation with a colleague and hold up its end of the conversation. It is

Actuaries: our value isn't email or reports or even models. It's **judgement, trust and long-term thinking**. We need to protect that.

"AI doesn't know the truth." That simple idea should change how you use every AI tool.

very good at filling in the blanks when I give a partial explanation of something, it can identify a set of information that completes the rest of that explanation. I call that interpolating. It is poor at extrapolating and the further into the future it is asked to look, the worse it gets at the extrapolation. It does not do any of what I would think of as reasoning, but it will claim that it has. It does not know what is true and what is false, but it will give you an answer if you ask it to tell you the truth. It seems to be getting better at answering problems that include math, though I would not trust it to do any calculations for me. You need to identify the specific strengths and weaknesses for your application.

We need to learn how to collaborate with AI in a way that aligns AI's strengths with our weaknesses and our strengths with its blind spots. I do that mostly in small steps. I rarely ask AI for an answer to my problem. Usually, I ask it for a reaction to my latest thought on the topic or I might share my reasoning and ask what is missing or what is another way to reach the same end. With small steps, I think that I am noticing its flights of fancy and can correct or redirect it. Usually that works, but sometimes, I will need to just stop and if I really need to make progress, I will shift my inquiry to a different model. Quite often, I do not necessarily use what it suggests, but I find that its answers might trigger me to think of a different path that I might not have found without that trigger. And I particularly like how I can stop the conversation at any time and pick it up an hour, a day or a week later right where we left off or I can clear the decks and start the conversation over from a blank slate. The strength that I always try to bring into these collaborations is my storehouse of real-world experiences. The AI will beat me at logic, not yet necessarily now, but someday soon. I need to bring intuition, my sense of right and wrong, my empathy for human situations and my ability to just know when the solution involves going off the path. My humanity.

We need to learn how and when we need to review and check and validate anything and everything that AI tells us, no matter how impressive it sounds. The phrase that I use to remind myself of that is that "AI doesn't know the truth." When it hallucinates, AI is not trying to pull one over on you, it is following the same logic that it always follows and this time, the sentence it makes happens to be something that you know is false. So, the "when" of checking on an AI response is almost all the time, potentially. Practically, the "when" depends on how you will be using the response and how much you know about its ability to give good answers in the domain it is operating within. I call this "Actuary in the Loop," and an actuary is needed when there is a strong need for broad domain expertise related to uncertainty and long-term financial security. The actuary in the loop will have that domain knowledge for insurance- and pension-related applications, along with the problem-solving skills to track down and help with repairing problems all guided by strong professional standards.

The actuary who is armed with this story and acts on these ideas will be able to make a greater contribution in any future and especially in a future that we share with AI.

Dave Ingram, FSA, CERA is an elected member of the SOA Board and an editor of this newsletter.

**ACTUARIAL
INTELLIGENCE
BULLETIN**

Actuarial AI Links

Artificial Intelligence Research from the SOA Research Institute
<https://www.soa.org/research/topics/artificial-intelligence-topic-landing/>

Artificial intelligence & machine learning from the SOA Emerging Topics Community
<https://www.soa.org/digital-publishing-platform/emerging-topics/>

Artificial Intelligence Research Listserv from the SOA Research Institute
<https://soa.wufoo.com/forms/join-ai-research-listserv>

American Academy of Actuaries AI Page
<https://actuary.org/topic/artificial-intelligence/>

AI Tools for Actuaries
<https://aitools4actuaries.com/>

Actuarial AI Case Studies
<https://github.com/IAA-AITF/Actuarial-AI-Case-Studies>

AI-Insights Podcasts from SOA Research Institute
<https://getpluggedin.libsyn.com/>

AI for Actuaries from the IAA AI Task Force
<https://www.aiforactuaries.org/>

Actuarial AI Case Studies
<https://github.com/IAA-AITF/Actuarial-AI-Case-Studies>



ACTUARIAL INTELLIGENCE BULLETIN

Editorial Committee

Jon Forster, ASA
Dave Ingram, FSA
Ronald Poon Affat, FSA
Frank Quan, PhD
Darryl Wagner, FSA

Associate Editor
Jing Kai Ong, ASA

Thank you for reading the September 2025 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