ACTUARIAL INTELLIGENCE BULLETIN



May 2025

Al Videos Are Most Viewed on SOA's Professional Development Edge+	2
Actuarial Intelligence Roundtable – Part 1	3
How to Use AI to Study for Actuarial Exams	6
Zero to Al Governance: Establishing a Principles-Based Framework	7
The Rise of LLMs in Risk Management	10
Bayesian Logistic Regression	11
Al Insights for Actuaries: A Virtual Symposium	13
IAA Artificial Intelligence Summit	13
Actuarial AI Links	15
Artificial Intelligence at the SOA Health Meeting	15
Editorial Committee	16
About the Society of Actuaries Research Institute	17

Welcome to the May 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.



AI Videos Are Most Viewed on SOA's Professional Development Edge+ JON FORESTER

The Society of Actuaries (SOA) launched a new professional development product, <u>Professional Development Edge+</u> (PD Edge+), at the 2024 ImpACT Conference. Since then, AI videos have been the top five viewed videos on the platform.

See the top 5 most viewed AI videos on PD Edge+ below:

1. Introduction to Al

Jing Lang, FSA, FCA, MAAA, FLMI is an AI expert and provides an insightful overview of Artificial Intelligence (AI) and its various manifestations, from generative AI generating divers content to everyday applications like smart speakers and self-driving cars. It distinguishes between Artificial General Intelligence (AGI), which can perform tasks akin to humans, and Artificial Narrow Intelligence (ANI), which focuses on specific tasks. The narrative delves into key AI subfields like Machine Learning (ML), Deep Learning (DL), and Generative AI (GenAI), explaining their functionalities and differences.

2. Gen Al Uses in Actuarial Science

In this video, Jing Lang explores AI's current and future landscape in the insurance industry, focusing on its impact on actuarial work. She highlights the wide-ranging applications of AI in business, particularly in customer interactions and marketing activities, while also discussing its potential disruptions in the insurance sector. Despite the advancements in AI, Jing emphasizes that certain aspects of actuarial work, such as pricing and reserving, remain largely unaffected for now but encourage proactive integration of AI tools for future adaptability.

3. History of Al

Jing Lang provides a panoramic view of the history and evolution of Artificial Intelligence (AI), from its theoretical origins to contemporary breakthroughs in deep learning, autonomous "I use PD Edge+ as part of the onboarding process for new members of my team to help them understand the big picture behind the work we do daily. It is a great starting point to build a baseline of knowledge from subject matter experts in an easily accessible format." Scott Damery, ASA, MAAA

"After a stint in retirement I'm looking to rejoin the workforce and PD Edge+ offered me an easy and economical way to obtain my CE credits to become compliant. It allowed me flexibility to take professional development that interests me and my goals. I really enjoyed the short form video content and courses" Sam Vorderstasse, FSA, MAAA

vehicles, and natural language processing. It highlights key milestones, such as the Turing Test and the development of chatbots, such as ELIZA, while emphasizing the ethical considerations that come with AI advancement, including privacy, bias, and transparency.

4. Intro to Gen Al

Generative AI is a frontier in artificial intelligence that can create new data from existing data sets, with applications in risk modeling, product development, data analysis, and more. This video presented by Chris Smith, ASA, MAAA, introduces the fundamentals, technologies, and implications of generative AI for actuarial science and the ethical considerations and limitations that actuaries must address. Learn how to harness the power of generative AI ethically and innovatively in your actuarial practice.



5. Risk Identification

Artificial Intelligence can bring many benefits to the insurance business, but also many risks that need to be identified and managed. In this video, Dave Ingram FSA MAAA, a SOA Board member and host of the podcast Crossing Thin Ice, discusses some of the common applications and risks of AI for insurance operations and actuarial work. He also shares his top 10 AI risks and some tips for finding the right balance between overconfidence and fear when adopting AI.

In addition to AI videos, there are two new AI courses on PD Edge+:

- 1. Al Use Cases for Actuaries: This course provides an overview of the role and potential of Artificial Intelligence (AI) in the actuarial profession. It covers various aspects, including AI layers and generative AI, and its applications in documentation, coding, data analysis, and routine tasks. AI tools and automation trends are also discussed. The course reviews several specific use cases, including utilizing AI for documentation and a life insurance model use case.
- Risks, Regulation, and Ethics: This course explores the applications and risks of Artificial Intelligence (AI) in the insurance industry, emphasizing the importance of a comprehensive AI Use Policy that covers data governance, ethical guidelines, transparency, fairness, human oversight, and compliance. It highlights the top risks in actuarial work, including bias and discrimination, and discusses best practices to mitigate these issues. Additionally, it reviews AI frameworks, generative AI risk management, regulatory trends, and the five pillars of ethical AI, along with a research report on avoiding unfair bias in insurance applications of AI.

About PD Edge+

<u>PD Edge+</u> is a subscription-based product with high-quality and dynamic professional development content in one location for your professional development. <u>Learn more</u> about PD Edge+ or contact Amy Schutzenhofer at <u>ASchutzenhofer@soa.org</u> or Jon Forster at <u>iforster@soa.org</u>.

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

ACTUARIAL INTELLIGENCE BULLETIN

Actuarial Intelligence Roundtable – Part 1

We gathered a dozen readers around a virtual table to talk about Actuaries and AI.

Is artificial intelligence (AI) a friend or foe to the actuarial profession?

Rob Bacher (Dallas, US) Al is a powerful tool for actuaries, enhancing productivity and analytical capabilities. However, its rise demands new skills in model validation, governance, and ethical interpretation to safeguard against unintended risks.

Jon Forster (Philadelphia, US) AI has the potential to support the profession, and becoming AI-enhanced actuaries may help ensure its continued robustness and relevance.



Jing Lang (Toronto, Canada) For an actuary who's open to new tools, technologies, and more efficient ways of working, AI is a total BFF. But for someone unwilling to learn or evolve, it's going to feel like a threat. Frankly, there's not much of a choice anymore. You either adapt—or you get left behind.

Frances Green (New York, US) There are so many great use cases, among them advanced modeling techniques and improved risk factor identification, that traditional actuarial methods might overlook.

Kailan Shang (Kitchener, Canada) I see AI as a friend to the actuarial profession. It's not something to fear but to embrace. Rather than replacing us, AI can empower us to do our work better and more effectively. I also believe we, as actuaries, have a unique role in guiding AI's development, ensuring that models are accurate, transparent, and fair.

Jacky Ng (Hong Kong, China) Actuaries who embrace AI will replace those who don't, like what spreadsheets did a few decades ago. Some may love it, others may loath it. AI, just like spreadsheets are here to stay – so we better make AI a friend than a foe!

David Schraub (Northbrook, US) A friend - in the same way as the introduction of grid computing, personal computer or the abacus are older friends.

Josee Deroy (New York, US) Not only is AI a friend, but it will provide tools to upscale the actuary's knowledge by offering additional insights to develop better and more comprehensive solutions in tackling financial risks.

Bruce B Rosner (New York, US) I see both the threat and opportunity and don't really know which one will win the day. On the one hand, most of us should be able to use AI to multiply our productivity in modeling and other work. On the other hand, we are already a small profession, so we need to expand our scope to stay relevant.

Arthur da Silva (New York, US) Friend, if used properly. Foe, if used improperly. Would that be a frenemy then?

JianGang He (Toronto, Canada) AI is definitely a friend to actuaries. It enhances our productivity, helps a lot with information gathering and organization, and expands our skill sets—especially in areas like programming and automation. By using AI, actuaries can focus more on higher-level insights, communication, and decision-making.

Amanda Hug (New York, US) AI is absolutely a friend. Actuaries have a long history of integrating new technology into our work, and AI is the next wave that will elevate what we do. It can reduce the time we spend on data manipulation and modeling, giving us more capacity to focus on business insights and strategic solutions. As a respected profession with a strong code of conduct, actuaries also have an important voice in conversations about AI governance and ethical use.

What are the most promising applications of AI in actuarial science today?

Jeff Yang (Los Angeles, US) I work in actuarial exam prep, so knowing which models are good for math, logic, and reasoning is important. The latest reasoning models (e.g. o3) significantly outperform the legacy models (e.g. 40) in these areas — so much so that the older models typically aren't even included in benchmark comparisons.

Frances Green (New York, US) Risk assessment and predictive modeling; personalization rather than risk pooling; scenario-based modeling that will assist in anticipating event severity and managing volatility, particularly in areas like cyber risk.

Kailan Shang (Kitchener, Canada) In my opinion, AI is transforming the actuarial field, particularly in predictive modeling and automation. Machine learning enhances risk assessment by identifying patterns in complex datasets,



leading to more accurate underwriting and pricing. Automation streamlines routine tasks like data processing and coding, allowing actuaries to focus on strategic decision-making.

David Schraub (Northbrook, US) Coding with the help of an LLM.

Arthur da Silva (New York, US) Lots of productivity gains on tedious tasks such as compiling information, summarizing, translating, transcribing, etc. AI can generate tests for models, intelligently find holes in product design that can be taken advantage of by policyholder behavior and conducting market research.

JianGang He (Toronto, Canada) There isn't one noticeable AI application in actuarial science right now, it's more about combining AI with what actuaries already do well. We are seeing useful tools in areas like regulatory and product search. Predictive models are still very effective, especially in data analysis and developing actuarial/underwriting rules.

Darryl Wagner (New York, US) Taking on repeatable well-defined tasks. Acting as a connecting "agent" for sub-tasks. Performing "first draft" analysis to save time and broaden the idea base. Providing conversational query capabilities against data and analysis sets.

What is one piece of advice you would give a colleague who wants to start experimenting with AI?

Rob Bacher (Dallas, US) Start by exploring user-friendly AI tools that address specific challenges in your area. Focus on incremental learning, and don't hesitate to seek guidance or collaborate with others who have experience.

If you're not sure how to start experimenting with AI, ask an AI model, "how do I start experimenting with AI?" Jeff Yang (Los Angeles, US) Jon Forster (Philadelphia, US) The most important step is to begin. The SOA provides useful content on AI through its meetings and PD Edge+ offerings that are suitable for beginners. It is recommended to start by understanding the available resources, then spend some time using an AI tool and experimenting with it.

Josee Deroy (New York, US) Solve one issue and expand from that point on.

Jing Lang (Toronto, Canada) Stop wanting. Start doing. Open a tool. Try a prompt. Break stuff. That's how you learn.

Frances Green (New York, US) Go for it! There is no better way to become comfortable with AI than with utilizing AI tools that your research and reasonable investigation has assured you are safe, ethical and trustworthy.

Kailan Shang (Kitchener, Canada) I would advise them to start by understanding Al's limitations so they can focus on what truly works today. Instead of chasing hype, experiment with AI in areas where it reliably adds value like automating repetitive tasks, enhancing data analysis or improving decision support.

Jacky Ng (Hong Kong, China) Al is a fast-moving beast – be agile, prepare to fail fast and move on. Don't waste your time brute forcing, as there are simply too many unknown pitfalls – reach out and seek help early when you are stuck and don't get frustrated – patience and perseverance is the key!

David Schraub (Northbrook, US) Just start.

Arthur da Silva (New York, US) Any time you find yourself frustrated with or dreading a manual task, try to see if AI can help a piece of that task, making sure you don't release sensitive data. Break down large complex tasks into



simpler components when doing this. If it doesn't work today, try again in a few months since AI tools are evolving so rapidly.

JianGang He (Toronto, Canada) My advice is simple: go and try—that's the only way to learn and to find how AI can support your work.

Darryl Wagner (New York, US) Don't wait. Jump in. Start small and in any direction – doesn't have to be actuarial.

ACTUARIAL INTELLIGENCE BULLETIN

How to Use AI to Study for Actuarial Exams

JEFF YANG

Significant advancements in generative artificial intelligence (AI) have allowed for the integration of powerful realtime tools when studying for actuarial exams.

Al offers numerous advantages over traditional learning methods:

- 1. It is always available, responds immediately, and is accessible globally.
- 2. It is patient, explains concepts well, and can provide a personalized experience.
- 3. It is constantly improving and can be readily trained on relevant subject matters.

The simplest use case of AI in actuarial exam prep is through AI chatbots that can explain concepts and help solve problems. Anyone can start using AI chatbots for free through popular chat-based apps such as ChatGPT. Many leading AI models are more than capable of answering basic questions, though their effectiveness can vary significantly depending on the exam and the material they were trained on.

For a more tailored experience, AI can be integrated directly into traditional exam prep tools and workflows. Examples include:

- 1. Automated grading for written exams, which is useful during practice and as a secondary review tool for actual exam graders.
- 2. Al-generated questions and solutions.
- 3. Customized AI models that incorporate relevant exam formulas and principles.
- 4. Learning recommendations based on a student's current progress.
- 5. A study buddy or moral support tutor for a less isolating experience.
- 6. Integration of AI-generated visuals, video, and audio.
- 7. Translation tools for international candidates.

One of the biggest challenges in using AI for actuarial exam prep is mitigating the risk of hallucinations (i.e., when AI responds confidently with factually incorrect statements). Several steps can be taken to mitigate the risk of hallucinations.

When asking about math calculations, candidates should prioritize using reasoning models.



- 1. When learning through practice questions, candidates often have access to the practice questions and solutions. In this case, AI can be used as an interpreter rather than a solutions engine. Providing the solution in the initial prompt and then asking AI to explain the solution will considerably reduce the likelihood of hallucinations.
- 2. When asking about math calculations, candidates should prioritize using reasoning models. Reasoning models, such as o3-mini, DeepSeek R1, Claude 3.7 Sonnet with reasoning enabled, or Gemini 2.5 Pro, "think" through a problem and are significantly more capable on actuarial exam questions that involve calculations. For example, non-reasoning models such as GPT-40 were able to answer 85% of the 600+ Exam P SOA sample questions correctly, whereas reasoning models consistently answered over 98% correctly. This difference is even more pronounced for Exam FM and FAM questions, although the absolute percentages are lower. Some leading AI model providers have also indicated that future models will no longer require users to choose between models. The default will be the best model for the task, which most likely will include reasoning capabilities if math is required.
- 3. Relying solely on AI tools for exam preparation is a risky approach. A more effective strategy is to supplement traditional study methods with AI. AI can be a helpful first responder when you need an immediate answer to a simple question. However, if the response seems questionable or incomplete, it can be helpful to verify with another trusted study source.

Al can be a powerful tool when preparing for actuarial exams. Whether you're seeking an explanation, looking for feedback on a solution, or just want a second opinion, Al can help. Just be mindful of potential hallucinations and always verify critical information.

Jeff Yang, FSA, is the founder of The Actuarial Nexus.

ACTUARIAL INTELLIGENCE BULLETIN

Zero to AI Governance: Establishing a Principles-Based Framework

I worked for a manager who was concerned that our governance was too hard to understand. We were at a large insurance company going through optimization. We needed to re-think our processes. He asked us to articulate key principles to guide our governance activities. Then, he asked us to summarize requirements to accompany each principle. Last, he had us articulate best practices to enable efficiency and consistency. He called this approach principles-based governance. This framework has helped me articulate the governance, and responsible and ethical use, of artificial intelligence in the insurance industry.

Agencies and organizations like the <u>World Economic Forum</u>, <u>European Union</u>, <u>White House</u>¹, <u>Federal Housing</u> <u>Finance Authority</u>, and <u>National Association of Insurance Commissioners</u> have articulated AI ethical principles. So have companies like <u>Google</u>², <u>Amazon</u>, and <u>Microsoft</u>.³ Using this representative sample set, we can input the ethical principles in a <u>word cloud generator</u>⁴ to see frequently represented words, below.





We can cross-reference the detailed principles to articulate a common set of ethical principles applicable to the insurance industry as follows:

- Avoid bias, both in data used to train a model and model output which may violate protected class requirements by law or regulation.
- Make results and uses of AI transparent. This includes adequate documentation of calculations, data, and training frequency, and making it clear to users they are interacting with an AI model.
- **Protect privacy and confidential information**, so we do not inadvertently feed it into an AI model or make it available as output.
- Ensure there is accountability, by appointing an accountable individual for each AI model.
- Ensure tools are reliable and safe, by ensuring AI models function as intended, avoid drifting, and remain stable and available.

Next, we can summarize the requirements for each principle. We can reference content like the <u>NAIC Model Bulletin</u> on the Use of AI Systems by Insurers, Ethical & Responsible Use of Data & Predictive Models - Executive Track, and the <u>FHFA Advisory Bulletin on AI/ML Risk Management</u>.

We can summarize representative requirements as follows:

- To **avoid bias**, we need data reviewed by a qualified individual to ensure it reflects the full range of users and conditions. We need a compliance review to ensure we do not violate protected class rules. We need ongoing output monitoring, including performance metrics, to ensure we introduce no new bias, especially for models re-trained frequently.
- To make results and uses of AI transparent, we need to register each AI model or tool in an inventory with its intended use. We need disclosure language, letting customers know they are interacting with an AI model and can speak with a human. We must ensure there is risk-based testing, documentation that addresses how the model was trained and on what data, and model monitoring so our model risk, audit, and regulatory stakeholders can evaluate the model.
- To protect privacy and confidential information, we must ensure we can disclose, upon request, how we use and store personal information. We must ensure we evaluate input data for personal and confidential information so we can flag it. We must ensure there is a method to detect personal or confidential information in model output to avoid inadvertent disclosure.
- To **ensure there is accountability**, we must assign an accountable party to each AI model or tool. That individual must attest periodically he or she understands requirements and is meeting them. We must assess whether we need a human review of output, and if so, incorporate this review for each use.



• To make sure tools are reliable and safe, we need ongoing performance monitoring for each model. We need reporting to ensure adequate performance and escalation protocols for non-performance. We must ensure previous versions are available if models become unreliable. We need to evaluate boundary scenarios to ensure the model can manage unlikely scenarios.

Last, we can articulate a set of best practices to enable us to meet requirements efficiently and consistently. Those are as follows:

- To **avoid bias**, we should document input data fields and catalog each element. We should include a peer review of training data to assess implicit bias.
- To make results and uses of Al transparent, we should create a decision tree, so owners understand what risk framework to follow model, operational, or other and register the tool accordingly. We should make points of contact (POCs), well versed in requirements, available for questions upon demand.
- To **protect privacy and confidential information**, we should provide training so model owners can identify it, know actions to take when they do, and so they understand to limit the use of this data unless necessary. We should build tools to check for personal or confidential information, to anonymize such data, and to filter it out of model output.
- To **ensure accountability**, we should provide periodic training for model owners to re-emphasize requirements and create a clear and easy checklist, so they understand them.
- To make sure tools are reliable and safe, we should create templates to store monitoring results so we can aggregate and report on them. We should have a periodic attestation by the model owner that the model is performing as intended. We should ensure plans for reverting to prior versions are well articulated, including when human override is necessary.

Bringing It All Together

We can now articulate principles, standard requirements, and best practices for governing the development, deployment, and use of AI. It looks like this:

Principle	Standard Requirements ⁵	Best Practices
Avoid bias	Review for representative data.	 Development data disclosure.
	Compliance review.	 Definition of data elements.
	 Ongoing monitoring/metrics. 	 Peer review of data.
Make results and uses of AI	 Inventory models/uses. 	Decision tree.
transparent	Disclosure language.	 Points of contact to support.
	• Document, test, and monitor.	
Protect privacy	• Disclosure personal info. (PI).	• Training to identify/minimize PI.
	Cross reference data for PI.	 Identify points of contact.
	• Ensure no PI in output.	 Tools to flag/anonymize PI.
Ensure there is accountability	 Assign an accountable party. 	Periodic training.
	 Periodic attestation. 	Requirements checklist.
	 May include human review. 	 Subject matter experts.
Make sure tools are reliable and	Monitor performance.	Output templates.
safe	 Establish reporting protocols. 	Periodic attestation.
	• Previous versions available.	Plans for prior versions.
	 Test boundary scenarios. 	

With the use of links — for example to a decision tree, output template, and specific requirements for documentation, testing, and monitoring — we can summarize an entire framework in one concise table.⁶ While this perspective lacks complexity and specificity —which are necessary to govern highly complex systems models — it can act as a starting point for employees to understand what is required, which best practices help to meet



requirements consistent and optimally, and which principles should guide their activities and encourage the responsible and ethical use of AI.

Endnotes

1: The Biden White House put out <u>A Blueprint for an AI Bill of Rights</u>, which contained the principles, in 2022.

- 2: Google has updated the <u>principles</u> since the first version they put out in 2018.
- 3: Click the click for each organization to see the AI ethical principles it has articulated.

4: <u>https://www.freewordcloudgenerator.com</u>.

5: Additional considerations for NAIC Model Bulletin Use of AI Systems compliance include 1) committee and management structure and strategy, 2) separate framework for AI or leveraging existing standards, 3) Internal Audit review, 4) third-party requirements, and 5) compliance evidence.

6: The above framework is representative and not meant to be definitive or prescriptive. Each company should separately assess the specific needs of an AI governance framework and evaluate applicable external guidance and requirements.

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



The Rise of LLMs in Risk Management

SYED RAZA

Transforming Risk Modeling with LLMs

Despite their rapid advancements, LLMs are still in the early stages of adoption within actuarial and risk management routines. Their acceptance hinges on accessibility and transparency—two aspects that have historically limited their use. Open-source models like DeepSeek are making strides in demystifying LLM decision-making by showcasing logical reasoning. This transparency could accelerate industry adoption, although ethical and intellectual property concerns remain.

LLMs particularly shine in regulatory compliance and risk reporting. They can monitor regulatory changes across jurisdictions, interpret legal documents, and highlight compliance requirements—reducing manual effort and ensuring thorough coverage. Additionally, these models enhance cybersecurity by detecting phishing attacks, insider threats, and other anomalies in real time. In fraud detection, LLMs analyze transaction patterns and customer interactions to flag suspicious activity before it escalates.

Elevating Credit Risk and Scenario Analysis

Beyond operational risk, LLMs offer significant potential in credit risk assessment. By integrating alternative data sources - such as business news and social media trends - with traditional credit metrics, they provide a deeper understanding of borrower creditworthiness. This approach is particularly valuable for evaluating small businesses and entrepreneurs who may lack extensive credit histories.

Another key advantage is the ability to generate and analyze hypothetical risk scenarios. LLMs can simulate economic downturns, cyberattacks, or regulatory shifts, allowing organizations to identify vulnerabilities and refine response strategies. For instance, asset managers can model the effects of interest rate changes on portfolios, leading to more informed investment decisions.



Case Study: GenAl for Reinsurance Optimization

To mitigate ethical concerns associated with publicly available LLMs, many companies are developing proprietary models tailored to their needs. LLMs can be built using languages like Python, extending beyond mainstream options like ChatGPT, DeepSeek, or Claude.

A groundbreaking application of Generative AI and Reinforcement Learning in reinsurance optimization illustrates this potential (Source: arXiv:2501.06404).¹ Reinsurance optimization is crucial for insurers to manage risk exposure and maintain solvency, yet traditional approaches struggle with dynamic claim distributions and evolving market conditions.

A novel hybrid framework combines Generative Models, such as Variational Autoencoders (VAEs), with Reinforcement Learning (RL) using Proximal Policy Optimization (PPO). This system:

- * Generates synthetic claims data, including rare catastrophic events, to address data scarcity.
- * Adapts reinsurance parameters dynamically, optimizing surplus while minimizing ruin probability.
- * Outperforms traditional models in scalability, efficiency, and robustness, achieving higher financial surpluses under stress-tested conditions.

This framework paves the way for broader applications in catastrophe modeling, multi-line insurance operations, and strategic risk-sharing.

The Future of LLMs in Risk Management

Looking ahead, LLMs are poised to revolutionize risk management further. We anticipate their growing role in realtime risk monitoring, automated stress testing, and dynamic risk reporting. Soon, customized LLMs may play a central role in actuarial tasks such as reserve calculations, pricing, reinsurance optimization, and risk modeling reshaping the future of financial risk management.

Endnotes

1: A Hybrid Framework for Reinsurance Optimization: Integrating Generative Models and Reinforcement Learning Stella C. Dong and James R. Finlay

Syed Raza, FSA, MAAA Actuarial Consultant & Founder of Actuarial 360 Solutions is title at organization (Style = Author info).



Bayesian Logistic Regression

JASON REED

Al systems used in actuarial work often rely on classification models trained on historical data. Bayesian Logistic Regression is a foundational method behind many such systems. Understanding how it works—and where it can go wrong—is crucial for actuaries evaluating or auditing AI tools for fairness, accuracy, and bias.

Many important actuarial applications involve the prediction of likelihood of an event occurring, such as hospital readmission, incurring catastrophically high claims in a future year, or policyholder lapse. A probability model



provides nuanced insight into the final binary classification, allowing the actuary to calibrate (based on tolerance for false positives/false negatives) how high a probability justifies a prediction that the event occurs.

For classification applications, standard regression techniques are inappropriate. For example, traditional risk scoring models involve solving a linear regression equation for which the dependent variable is relative medical claim costs, which can be any number greater than or equal to zero. On the other hand, the tools of Logistic Regression allow us to predict the probability of an event as a function of predictive variables (like age, gender, prior claims, disease status, and interactions among these predictors).

The Difficulty of Predicting Rare Events

Logistic regression works best when classes are close to balanced (there is a significant proportion of training set data in both classes). Most actuarial applications involve predicting rare events, where classes are (usually very) imbalanced. For example, predicting catastrophically high-cost claimants or mortality in the working age population. While maximum likelihood estimators for logistic regression are unbiased when both classes have sufficient representation, it becomes biased when one class is greatly over-represented (and therefore fit well with high accuracy) and the other is under-represented. Why? The under-represented class is more likely to be missing outlying observations multiple standard deviations from the mean, causing estimates of the distribution to be under-weighted in the tails. This leads to expected error (bias) in the direction of underpredicting the probability of the rare event class. Unfortunately, that error would lead to greater exposure to losses from these rare events (i.e., underpricing due to underestimation of the likelihood of the rare event). Fortunately, there are known computational ways to estimate and correct for this bias, but it is critical not to neglect this potential prediction error. For example, resampling technique, custom loss function, and ensemble, etc.¹

Bayes Theorem and the Bayesian Statistical Perspective

The Bayesian way of thinking about statistics relies on the idea that we have prior information which we bring to our statistical model, not only the data we collect in the sample. This prior information generally applies to the parameters of the model we are fitting. Then after collecting a sample (assumed to be from a different distribution), we update our estimates of these parameters, reflecting a blend of prior information and actual collected data. With that posterior distribution, we can make statistical inferences.

Why Bayesian Logistic Regression?

The Bayesian paradigm allows us to reflect additional uncertainty in our model. For example, a given member's risk score is 1.2 (expected costs 20% higher than average), but we may want to add uncertainty to this best estimate prediction (i.e. there is a high likelihood that costs will be between 15% and 25% greater than average). The Bayesian paradigm also allows us to incorporate past knowledge or expert opinion about the parameters we are estimating, which can come from sources outside of the sample data we have collected. Alternatively, a Bayesian "uninformative prior" can allow us to reflect variance in the parameter without presuming that some values are more likely than others.

Bayesian analysis is also valuable because it combines prior expert knowledge (actuarial assumptions) with actual sample data collected to produce a posterior distribution of likely outcomes, yielding a blend of prior and databased information. This posterior distribution enables quantification of uncertainty in parameter estimates, understanding of variance and skewness in the distribution of outcomes, and likelihood of extreme events. Bayesian models are also dynamic with respect to the collection of additional data. This feature allows us to continuously update our models with new information as it comes in, rather than starting over each time we want to conduct an experience study.



In the risk score models considered above, we can reflect variability in the contribution to risk measured by each predictor. That variability can be measured by just one parameter (i.e., standard deviation), or we can allow flexibility in the distribution of the model parameters. Bayesian modeling can allow for very flexible models, with the caveat that the actuary must have enough data to measure and model these quantities appropriately.

Endnotes

1[:] Hu, C., Quan, Z., & Chong, W. F. (2022). Imbalanced learning for insurance using modified loss functions in tree-based models. *Insurance: Mathematics and Economics*, *106*, 13-32.

Jason Reed, FSA, MS, MAAA Senior Director, UnitedHealthCare Community and State Actuarial



AI Insights for Actuaries: A Virtual Symposium

Following 2024's inaugural and successful AI Insights for Actuaries: A Virtual Symposium, we have an exciting line-up for 2025. Over two days, attendees will explore AI's transformative impact on actuarial science through a practical and a governance/management track. The symposium offers insights into emerging AI technologies, practical applications, and the future of actuarial work globally, helping actuaries thrive in a changing environment. Plus, don't miss the sessions on professionalism and bias. *AI Insights for Actuaries: A Virtual Symposium | SOA*

Join us for the free opening session of AI Insights for Actuaries: A Virtual Symposium. This complimentary session will explore the fascinating world of artificial intelligence (AI), providing foundational knowledge on AI's trajectory and practical insights on how actuaries can harness AI to enhance their daily work. Designed for all audiences and skill levels, this session promises to deliver cutting-edge information that is both accessible and impactful. *AI Insights for Actuaries: The Essential Overview*



IAA Artificial Intelligence Summit

JACKY NG

San Francisco, the home of Silicon Valley and innovations. It is a rather inspirational location to host the second IAA AI Summit with the theme of **The AI Enabled Actuary**. The Summit was held on 20th – 21st Feb 2025 and was well attended by over 50 delegates representing 20 full member associations (FMAs) from around the globe.

Charles Cowling, past president IAA opened the proceedings by sharing the missions of IAA are:

- 1. Impact global stakeholders
- 2. Assure the reputation of the profession
- 3. Advance the competency of the profession



With the backdrop of advances in artificial intelligence, how can the actuaries from all corners of the world seize the opportunities in propelling the profession forward and for actuaries to remain relevant?

The message from Charles was quite clear. Actuaries need to embrace AI. There is absolutely no doubt actuaries with AI will replace those actuaries without AI. There is no time to waste. We need to seize the moment. The momentum is already there, with Charles reflecting on the inaugural IAA AI Summit held in Singapore back in April 2024. In the short span of eight months that followed, the AI Taskforce went full steam ahead in meeting the deliverables deadline by the end of 2024. There were five workstreams, each with their own set of deliverables:

- 1. Professionalism and Ethics Developing guidelines and principles for actuaries
- 2. Education Developing enhanced curricula and training modules to include AI
- 3. Changing role of Actuaries Developing basic definitions of AI; areas where AI can augment actuarial work; encouraging the development of appropriate actuarial expertise
- 4. Governance Monitoring and evaluating governance frameworks, policies, and regulations; engaging with regulators, standard-setting bodies, and policymakers
- 5. Innovation Cultivating a growth mindset among actuaries

The workstream leads presented their respective achievements. One particular deliverable from the Innovation workstream was the collaboration platform https://www.AlforActuaries.org/ which is now live! Deliverables from the IAA AI Taskforce will be made available on the platform in due course.

We experienced one of the liveliest interactive Q&A session, as delegates contemplated what are holding actuaries back from embracing AI, or was it simply the wrong perceptions projected from the outside world? Certainly, there are challenges that the actuarial profession need to overcome. One question that sparked a lively debate was how to evaluate the performance of AI or LLM models as some of the traditional evaluation techniques may no longer be applicable – so can the AI models be really trusted over traditional actuarial or statistical models?

These are healthy debates the actuarial profession needs to have as actuaries from around the world chart their own path in evaluating and embedding AI in their workflow. The IAA AI Taskforce, with support from the FMAs, including the SOA, is certainly committed to making this journey a success. To this end, having done the groundwork during 2024 on the "principles", the AI Taskforce will turn its focus to "practice" in equipping the AI Enable Actuary. The workstreams have been consolidated down to four with the term lengthen to two years until 2026. SOA is represented by a delegate in each of the workstream:

- 1. Engagement and Foundations Promoting engagement, building and growing the community of actuaries using AI, as well as providing an easy way for actuaries to get started in AI. (SOA delegate: Robert Eaton)
- 2. Research and Advancement Scanning the AI environment relevant to actuaries and supporting the education of actuaries in this field. (SOA delegate: Toby Hall)
- 3. Case Studies and Tools Collect, house, and organize tools and workbenches for actuaries to use in practice (SOA delegate: Jacky Ng)
- 4. Adoption Framework Focus on adoption, incorporating best practices, professionalism, governance, ethics as well as practicality aspects. (SOA delegate: Dorothy Andrews)

Each of the four workstreams had breakout sessions during the Summit to explore their focus areas and setting up the structure for next steps, which is now set in motion. Stay tuned for more updates as we progress through the journey of The AI Enabled Actuary.

Jacky Ng Data Science Director APAC, Director, Group Data Services Swiss Re Asia Pte. Ltd.





Actuarial AI Links

Artificial Intelligence Research from the SOA Research Institute

Artificial intelligence & machine learning from the SOA Emerging Topics Community

Artificial Intelligence Research Listserv from the SOA Research Institute

Al for Actuaries from the IAA AI Task Force

ACTUARIAL INTELLIGENCE BULLETIN

Artificial Intelligence at the SOA Health Meeting

DAVID STODDARD

The annual SOA Health Meeting will take place both in-person and virtually. The in-person meeting will take place in Dallas, TX from June 22 - 25 and will offer up to 15 hours of continuing education credits, while the virtual meeting will take place July 22 - 24 and offer up to 10 hours of continuing education credits. As part of both agendas, there are quite a few sessions that will focus on AI. We have identified the following list and would recommend anyone interested in learning more about how AI can be used in the actuarial profession to attend some or all these sessions.2025 SOA Health Meeting, June 22-25 in Dallas, TX https://www.soa.org/prof-dev/events/2025-health/

Session 1F: The Math Behind the Machine - A Review of First Principles Session 2F: Artificial Intelligence and Actuarial Modeling - Ask an Expert! Session 7B: Health AI in Practice - Lessons from Top U.S. Health Systems Session 9C: Adventures in AI - Group Insurance Perspectives Session 9F: GenAI in Action - Elevating Health Actuaries to New Heights Session 9G: AI and ASOPs - Re-examining Standards in Practice Session 10F: Debating AI - Valuable Tool or Overhyped Trend?

Virtual 2025 SOA Health Meeting, July 22-24 https://www.soa.org/prof-dev/events/2025-health-virtual/

Virtual - Session 1A: From Actuary to Data Scientist - Career Journeys into Al Virtual - Session 2A: Building and Using an Al Simulation Model Virtual - Session 3A: But We have ChatGPT at home... Or Operationalizing Generative Al

David Stoddard FSA Director Health and Benefits, WTW



ACTUARIAL INTELLIGENCE BULLETIN

In recognition of the rapidly evolving nature of artificial intelligence, this publication embraces a diversity of perspectives. The views expressed in individual articles are those of the authors and may not reflect the consensus of the broader AI community—or even of the editorial committee. Some opinions may represent minority or emerging viewpoints. However, it is the firm position of the editorial committee that the free and uncensored exchange of ideas is essential for progress. Accordingly, all contributions are welcomed, provided that statements are not blatantly incorrect or misleading.

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 May 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, datadriven 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 <u>strategic research programs</u>: 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 <u>topical research available</u>, 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 <u>www.SOA.org</u>