The Actuary in an Age of AI

Presented by **Ronald Richman FIA FASSA CERA CPCU**

CEO and Founder, InsureAI

ron@insurai.co

Talk based on

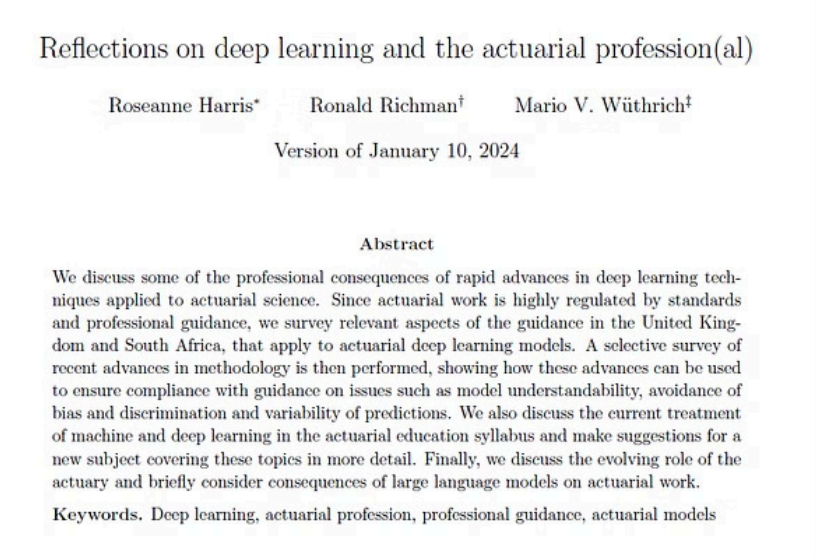

Reflections on deep learning and the actuarial profession(al)

Roseanne Harris* Ronald Richman† Mario V. Wüthrich‡

Version of January 10, 2024

Abstract

We discuss some of the professional consequences of rapid advances in deep learning techniques applied to actuarial science. Since actuarial work is highly regulated by standards and professional guidance, we survey relevant aspects of the guidance in the United Kingdom and South Africa, that apply to actuarial deep learning models. A selective survey of recent advances in methodology is then performed, showing how these advances can be used to ensure compliance with guidance on issues such as model understandability, avoidance of bias and discrimination and variability of predictions. We also discuss the current treatment of machine and deep learning in the actuarial education syllabus and make suggestions for a new subject covering these topics in more detail. Finally, we discuss the evolving role of the actuary and briefly consider consequences of large language models on actuarial work.

Keywords. Deep learning, actuarial profession, professional guidance, actuarial models

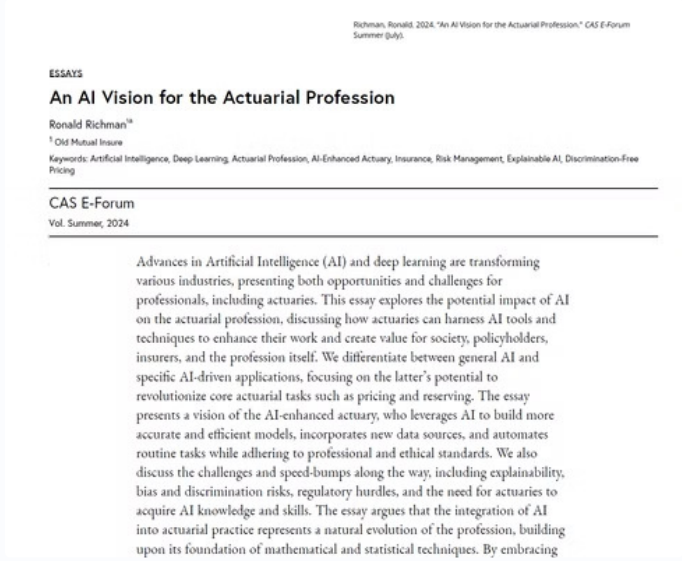

Richman, Ronald. 2024. "An AI Vision for the Actuarial Profession." *CAS E-Forum* Summer (July).

ESSAYS

An AI Vision for the Actuarial Profession

Ronald Richman[1a]

[1] Old Mutual Insure

Keywords: Artificial Intelligence, Deep Learning, Actuarial Profession, AI-Enhanced Actuary, Insurance, Risk Management, Explainable AI, Discrimination-Free Pricing

CAS E-Forum

Vol. Summer, 2024

Advances in Artificial Intelligence (AI) and deep learning are transforming various industries, presenting both opportunities and challenges for professionals, including actuaries. This essay explores the potential impact of AI on the actuarial profession, discussing how actuaries can harness AI tools and techniques to enhance their work and create value for society, policyholders, insurers, and the profession itself. We differentiate between general AI and specific AI-driven applications, focusing on the latter's potential to revolutionize core actuarial tasks such as pricing and reserving. The essay presents a vision of the AI-enhanced actuary, who leverages AI to build more accurate and efficient models, incorporates new data sources, and automates routine tasks while adhering to professional and ethical standards. We also discuss the challenges and speed-bumps along the way, including explainability, bias and discrimination risks, regulatory hurdles, and the need for actuaries to acquire AI knowledge and skills. The essay argues that the integration of AI into actuarial practice represents a natural evolution of the profession, building upon its foundation of mathematical and statistical techniques. By embracing

- **Reflections on deep learning and the actuarial profession(al) - https://eforum.casact.org/article/120560-an-ai-vision-for-the-actuarial-profession**
- **An AI Vision for the Actuarial Profession - https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4758296**
- **Credibility Transformer - https://arxiv.org/abs/2409.16653**
- **Deep Learning with Whittaker and Henderson - tbd**

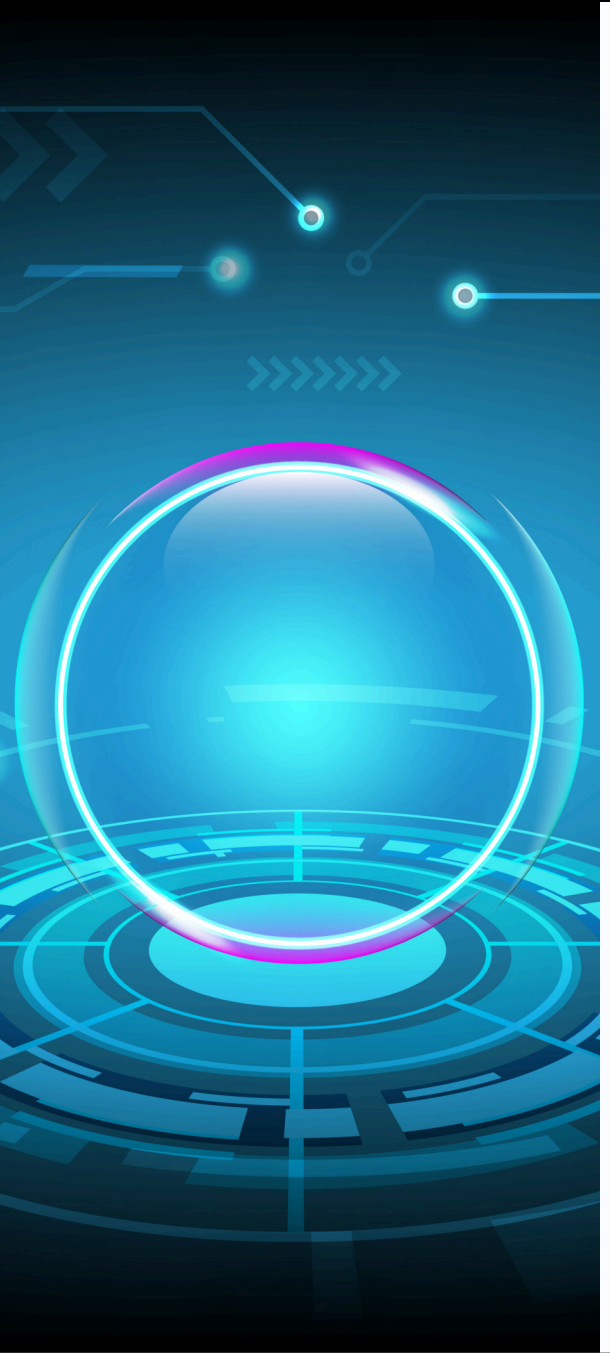

Introduction

1 Deep Learning Revolution

AI and deep learning are revolutionizing various industries, signaling a paradigm shift in how data is analyzed and decisions are made.

2 Actuarial Adaptation

The actuarial profession must adapt to remain competitive and relevant in an increasingly data-driven world.

3 Presentation Roadmap

This presentation will cover:

- Setting the stage
- The AI-enhanced actuary
- Actuarial AI applications
- Technical Challenges
- The path ahead

Situational Awareness

- Image is from "Situational Awareness" by Leopold Aschenbrenner
- https://situational-awareness.ai/
- Imagine what another order of magnitude jump in intelligence of GPT-type models could mean?
- Consistent improvements in compute/algorithmic efficiencies/unlocking latent capabilities
- "GPT-2 to GPT-4 took us from ~preschooler to ~smart highschooler; from barely being able to output a few cohesive sentences to acing high-school exams and being a useful coding assistant. That was an insane jump. If this is the intelligence gap we'll cover once more, where will that take us? … Likely, it will take us to models that can outperform PhDs and the best experts in a field."

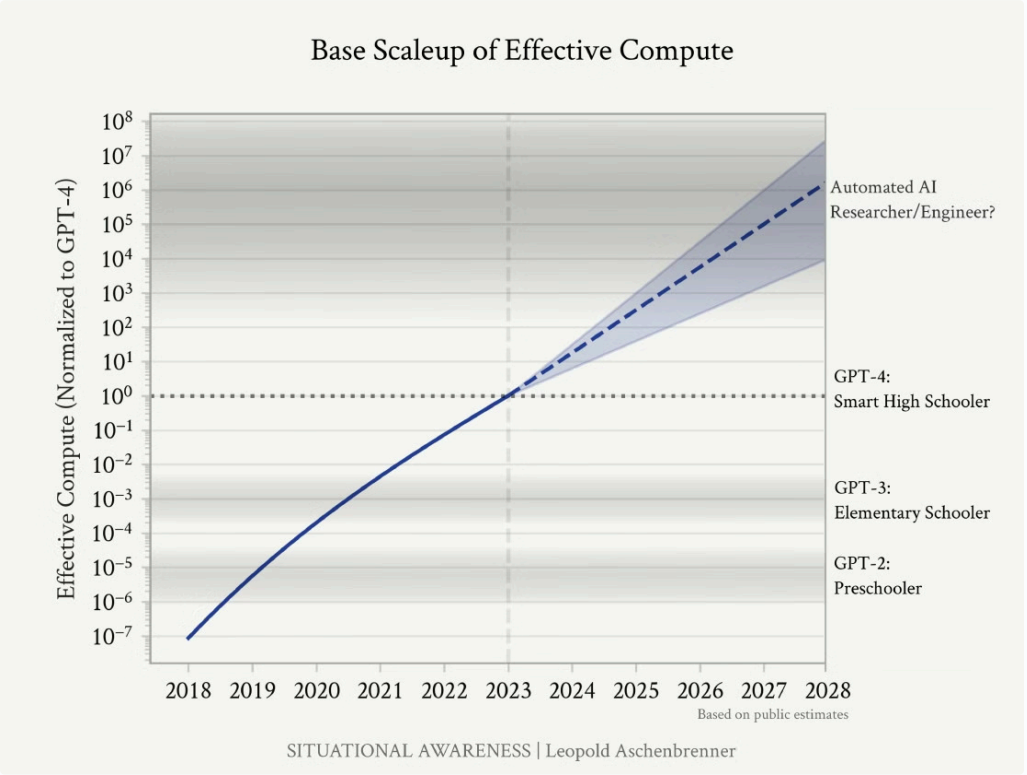

Situational Awareness for Actuaries

- Imagine an actuarial model powered by advanced AI that generates **personalized mortality tables** for every policyholder in a life insurance book.
- **Granularity and Precision**: Instead of broad cohorts, the model generates mortality predictions at an individual level, dynamically adjusting for factors like **lifestyle, genetic predispositions, socioeconomic data, and even real-time health monitoring**.
- **Dynamic Updates**: Mortality rates are no longer static. They evolve in real-time based on changes in individual health data, environmental risks, or advances in medical science.
- **Algorithmic Efficiencies**: Advanced deep learning reduces the computational burden, making it feasible to process massive datasets for millions of policyholders while remaining cost-effective.

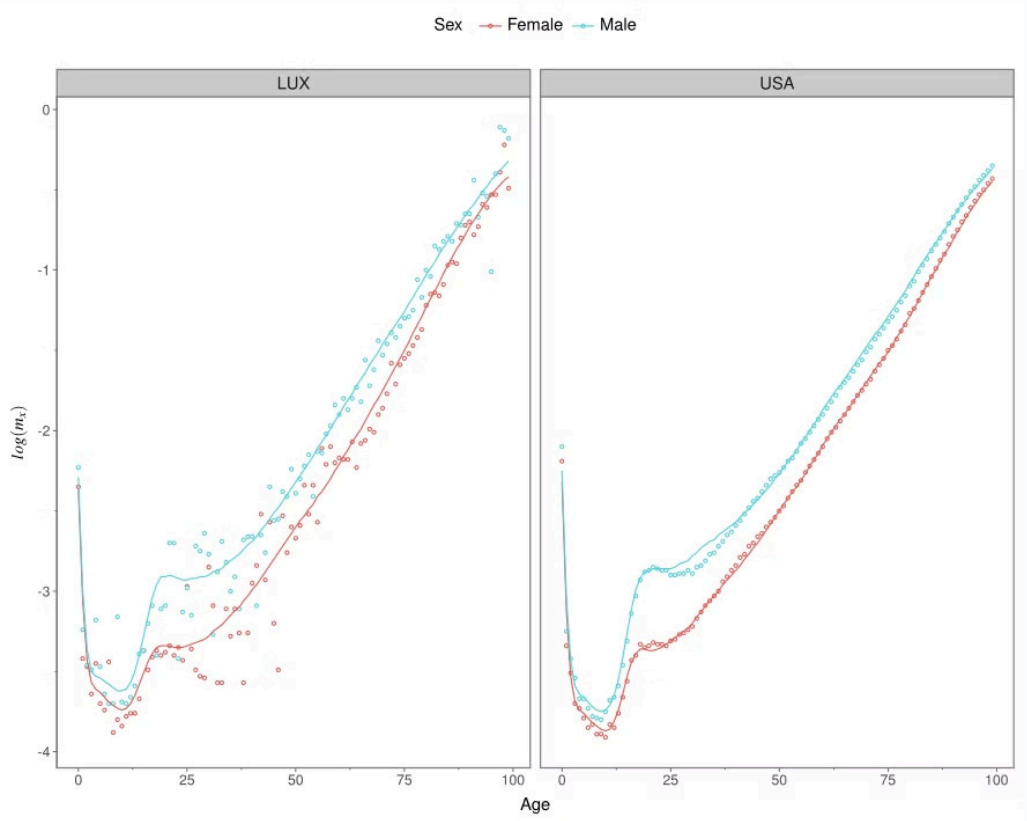

Just like we are starting to see a vision of what AGI could look like in a few years time, so too what could actuarial AI look like in 2040?

We are beginning to understand how artificial general intelligence (AGI) might emerge. This same exploration can guide us towards envisioning actuarial AI.

Narrow/Specific vs General AI (1)

Narrow AI

Narrow AI refers to AI-driven applications designed to perform specific tasks, such as pricing or reserving, which are at the core of the technical work actuaries perform. These tools are **tailored to solve particular actuarial problems** and are not necessarily applicable outside of the actuarial domain.

General AI

General AI, on the other hand, refers to AI systems that can potentially perform as well as, or better than, humans on a wide array of tasks. Large Language Models (LLMs) like GPT-4 and Claude 3 are (perhaps?) examples starting to approach a **more general form of AI**.

Narrow/Specific vs General AI (2)

Narrow vs General?

Specific AI tools being developed in actuarial literature may hold equal or more promise for the actuarial profession. These narrow AI tools are designed to revolutionize core actuarial tasks and can be directly integrated into actuarial work.

Think beyond the hype

While LLMs may be the first thing that comes to mind when thinking about AI, they are not the only aspect to consider. The integration of narrow AI into actuarial practice represents a natural evolution of the profession, building upon its foundation of mathematical and statistical techniques.

Why?

Narrow AI tools allow actuaries to build models more quickly and attain greater accuracy than traditional methodologies, incorporate new data sources directly into their models, and utilize new classes of inherently explainable models.

Goal

By embracing these specific AI tools and developing new skills, actuaries can unlock opportunities for innovation, efficiency, and value creation within the insurance industry and beyond.

The tools exist: (narrow) AI in Core Actuarial Tasks

1. **Pricing**

   AI enables more accurate and granular risk segmentation, leading to personalized pricing strategies that better reflect individual risk profiles.

2. **Reserving**

   Large-scale neural network models and automated assumption selection streamline the reserving process, enhancing accuracy and efficiency.

3. **Risk Management**

   Improved risk assessment and mitigation strategies emerge from the advanced analytical capabilities provided by AI.

See reading list at end of deck

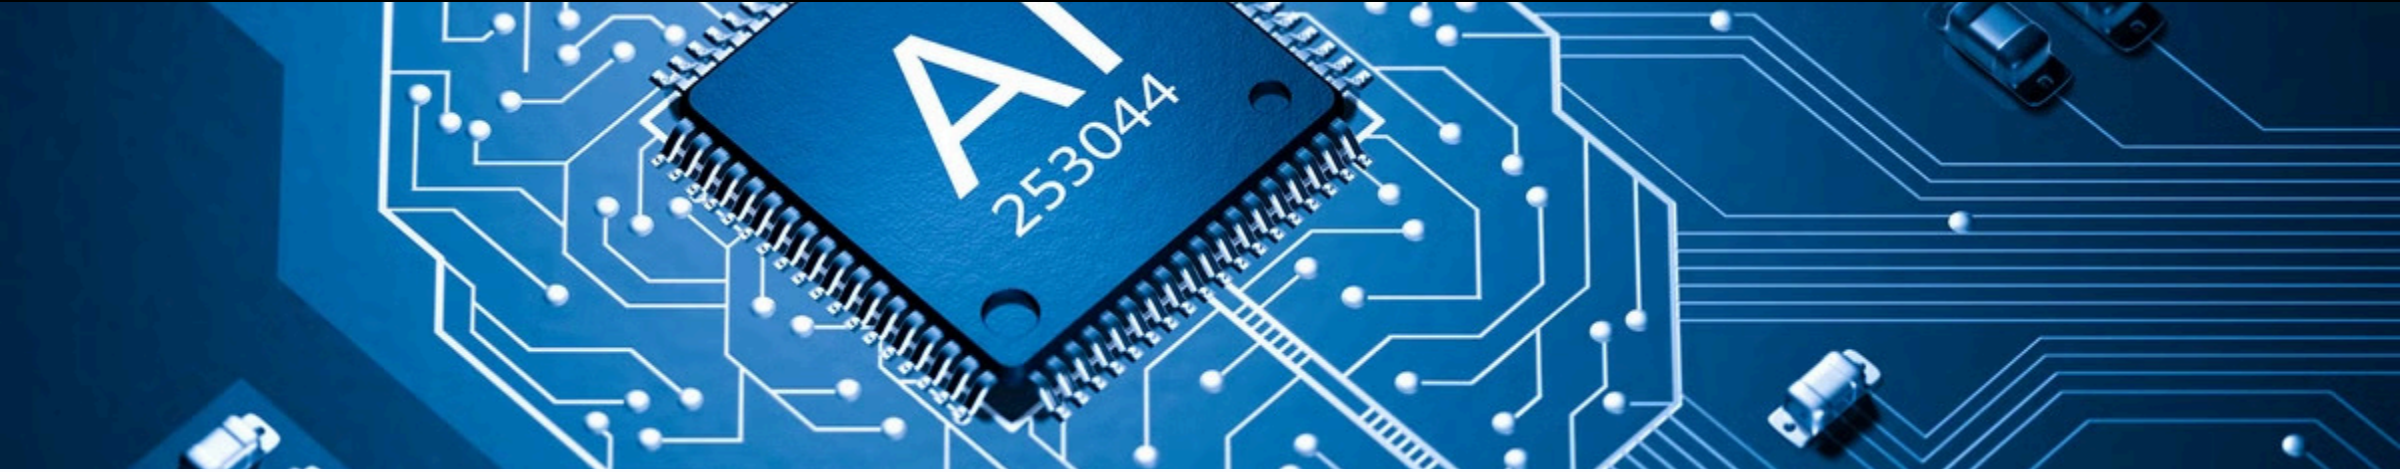

Benefits of AI in Actuarial Practice

Model Accuracy

AI contributes to increased accuracy and efficiency in model building, transforming the predictive capabilities of actuaries.

New Data Integration

The incorporation of new data sources, such as telematics and wearables, enriches the data landscape for actuaries.

Automation of Tasks

AI automates routine tasks, freeing actuaries to focus on higher-level strategic initiatives and complex problem-solving.

Actuaries need to "level up"

Evolution of AI Tools

AI tools and techniques, such as machine learning and deep learning, are already being developed and applied across various industries, including insurance and finance.

New Perspective

The actuarial profession has begun to explore the potential of AI in areas such as pricing, reserving, and risk management, with a growing body of research demonstrating the benefits of these approaches.

Slow pace of change

However, the adoption of AI in actuarial practice has been **relatively slow,** with many actuaries still relying on traditional methods and tools.

A new approach is needed

To fully harness the power of AI and remain competitive in an evolving landscape, the actuarial profession needs to **"level up"** and embrace a new perspective: the AI-enhanced actuary.

The AI-Enhanced Actuary

- **Empowering Actuarial Science with AI**: The AI-enhanced actuary combines the best of both worlds: the rigorous mathematical and statistical training of traditional actuarial science, and the cutting-edge predictive power and flexibility of AI techniques.
- **Mastering AI Tools and Techniques**: This new perspective involves acquiring a strong foundation in AI and machine learning concepts, as well as practical skills in implementing and interpreting AI models.
- **Adapting AI to Actuarial Work**: The AI-enhanced actuary is not just a user of AI tools, but a professional who deeply understands and can adapt these tools to meet the specific needs and requirements of actuarial work.

Unpacking the AI-Enhanced Actuary

Combines traditional actuarial expertise with cutting-edge AI knowledge and skills.

Leverages AI tools and techniques to enhance accuracy, efficiency, and value creation in actuarial tasks such as pricing, reserving, and risk management.

Adapts AI methods to comply with professional and ethical standards.

Ensuring transparency, fairness, and accountability.

Embraces continuous learning and upskilling.

To stay at the forefront of AI advancements and their applications in actuarial science.

Collaborates with professionals from other disciplines.

Such as data science and machine learning, to develop innovative solutions and drive AI integration in the insurance industry.

Applies AI to incorporate new data sources.

Automate routine tasks and generate insights that enable data-driven decision-making and strategic analysis.

Contributes to the responsible development and deployment of AI systems.

Leveraging actuarial expertise in risk assessment, governance, and ethical considerations.

Shapes the future of the actuarial profession.

The perspective of the AI-enhanced actuary is not just a matter of staying relevant in a changing world; it is an opportunity for the actuarial profession to redefine its value proposition and make an even greater impact in an AI-driven future.

Example: Current State of Reserving

Traditional Reserving Techniques

Many reserving analyses still rely on traditional methods like chain-ladder or Bornhuetter-Ferguson (BF).

The process involves manually comparing actual reserve development to expected (AvE) to assess accuracy and sufficiency.

Actuaries then manually select new loss development factors and assumptions for each line of business.

Manual and Time-Consuming Process

The current reserving process is largely manual, time-consuming, and relies heavily on individual actuaries' expertise.

Some less experienced actuaries may apply methods mechanically and miss out on worrying trends.

Common trends across lines of business are usually only discovered towards the end of the process.

Limited Integration of Advanced Techniques

There is limited integration of advanced modeling techniques, such as machine learning or AI, in the traditional reserving workflow.

The current approach may not fully capture the complex interactions and dependencies between different lines of business.

Peer Review and Communication

The analysis results in a final set of selected reserves, which are peer-reviewed and communicated to the business and senior management.

Discussions may lead to modifications of the reserves before final figures are booked into accounts and regulatory templates.

Writing the reserving report can take days of a skilled actuary's time.

Example: The Future of Reserving with AI

1 Automated Data Extraction

The journey begins with an automated data extraction and reconciliation process. An AI-driven agent reviews and queries significant data discrepancies.

2 Actual vs. Expected Analysis

Another AI system analyzes the actual vs. expected exhibits, identifying critical concerns and detecting common trends.

3 Neural Network Reserving Model

The reserving actuary updates a large-scale neural network reserving model, enhancing credibility and accuracy.

4 Automated Reserving Procedures

The initial reserves are compared to other automated reserving procedures, selecting the best combination of techniques.

5 Summary and Presentation

A Large Language Model summarizes the results and builds a presentation, describing key changes, insights, and financial impacts.

In this future state, AI enhances the reserving process by automating routine tasks, incorporating advanced modeling techniques, and enabling actuaries to focus on high-level analysis and communication. The result is a more efficient, accurate, and insightful reserving process that leverages the strengths of both AI and human expertise.

Real-world issues in an AI-Enhanced Reserving Process

New Line of Business

Scenario: During the reserving process, the actuarial team discovers that a new line of business, covering errors and omissions (E&O) and crisis management for AI-based systems, was launched in the past year.

Challenge: Limited data availability, with only three quarters of reserve development, and noisy loss experience due to small business volumes.

Solution:

- The actuary fine-tunes a small parameter vector for the new line of business using the main company reserving neural network.
- By comparing the fine-tuned parameters to other parameters in the model, the actuary identifies that the most similar lines of business are cyber insurance and product recall.
- Using this insight and expert judgment, the actuary recommends the indicated reserves for booking.

Importance: Adapting AI models to handle new lines of business with limited data is crucial for accurate reserving and strategic decision-making.

External Auditor Challenge

Scenario: During the annual reserve audit, external auditors challenge the reserving actuaries on the neural network model's allowance for significant negative development in several long-tailed lines of business.

Challenge: The market practice is to disallow negative development and include a tail factor instead.

Solution:

- The actuary investigates methods to align AI models with human feedback and discovers techniques such as incorporating prior judgments through Bayesian methods or imposing constraints on the network.
- After implementing these modifications, judgment-based tail factors are incorporated into selected lines of business.

Importance: Integrating human expertise and market practices with AI models ensures the reserving process is aligned with industry standards and stakeholder expectations.

Actuaries must be prepared to:

1. Leverage AI models to handle new or unique situations, such as limited data for new lines of business, by applying transfer learning techniques and incorporating expert judgment.
2. Respond to external challenges and align AI models with industry practices and stakeholder expectations by integrating human feedback and domain knowledge into the modeling process.

Adapting Deep Learning for Actuarial Purposes

Incorporating Actuarial Thinking

Recent approaches to applying deep learning to actuarial prediction tasks have focused on using traditional approaches within advanced models.

Model Explainability

Actuaries are adapting deep learning models to ensure they meet professional standards for explainability, using techniques like SHAP, PDPs, and ICE plots for post-hoc interpretability, and inherently interpretable models.

Mitigating Bias

To ensure fairness, actuaries use discrimination-free pricing techniques and multi-task networks to mitigate bias and comply with ethical standards.

Addressing Variability

Network aggregation, or nagging, is employed to address variability and reproducibility challenges in neural network predictions.

Actuaries can improve AI models with traditional thinking

- Recent works at the intersection of AI and actuarial science show how deep learning models enhanced with actuarial thinking can improve model performance!
- Example - the Credibility Transformer
- Credibility is the age-old actuarial practice of blending individual and portfolio experience
- Usually thought about when setting rates for one portfolio of risk among many others
- Can we apply this thinking in the cutting edge world of Transformers?

Attention Is All You Need

ish Vaswani* ogle Brain ıi@google.com

Noam Shazeer* Google Brain noam@google.com

Niki Parmar* Google Research nikip@google.com

Jakob Us Google Re usz@goog

lion Jones* gle Research ı@google.com

Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu

Łukasz Kaiser* Google Brain lukaszkaiser@google

Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Abstract

he dominant sequence transduction models are based on complex recurrent o onvolutional neural networks that include an encoder and a decoder. The bes erforming models also connect the encoder and decoder through an attentio echanism. We propose a new simple network architecture, the Transforme ased solely on attention mechanisms, dispensing with recurrence and convolution ntirely. Experiments on two machine translation tasks show these models t superior in quality while being more parallelizable and requiring significantl ss time to train. Our model achieves 28.4 BLEU on the WMT 2014 English -German translation task, improving over the existing best results, includin sembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task r model establishes a new single-model state-of-the-art BLEU score of 41.8 afte aining for 3.5 days on eight GPUs, a small fraction of the training costs of th est models from the literature. We show that the Transformer generalizes well t her tasks by applying it successfully to English constituency parsing both wit rge and limited training data.

Credibility Transformer

- We can introduce credibility principles into how we train AI models.
- Instead of only relying on the data we see, ensure that the model calibrates to what we know about the portfolio in general
- Automatically assign credibility between individual covariates and portfolio specific information
- Improves model performance even within a state-of-the-art Transformer setup (for tabular data)

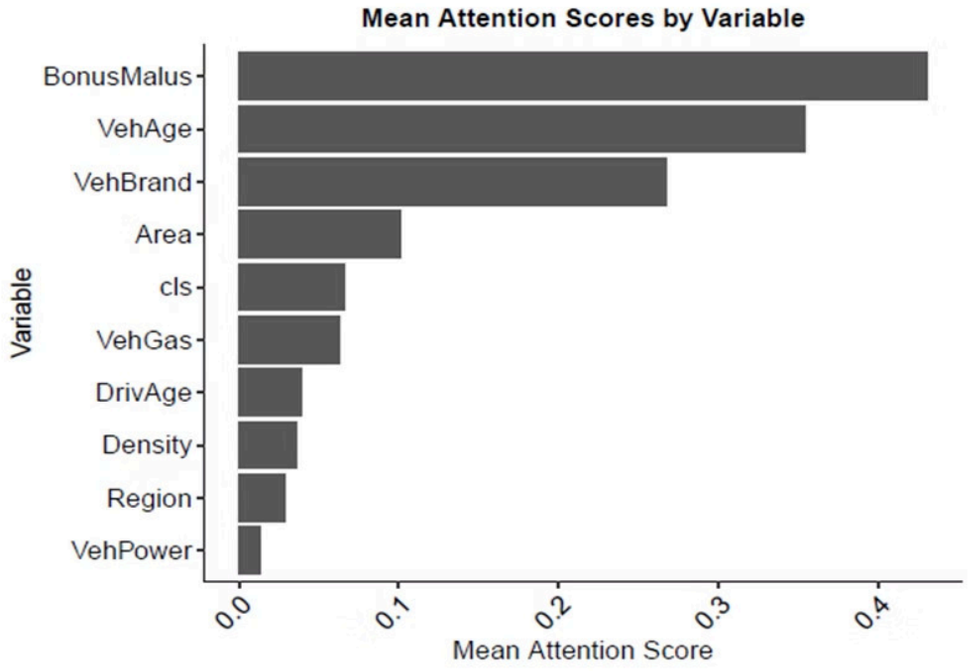

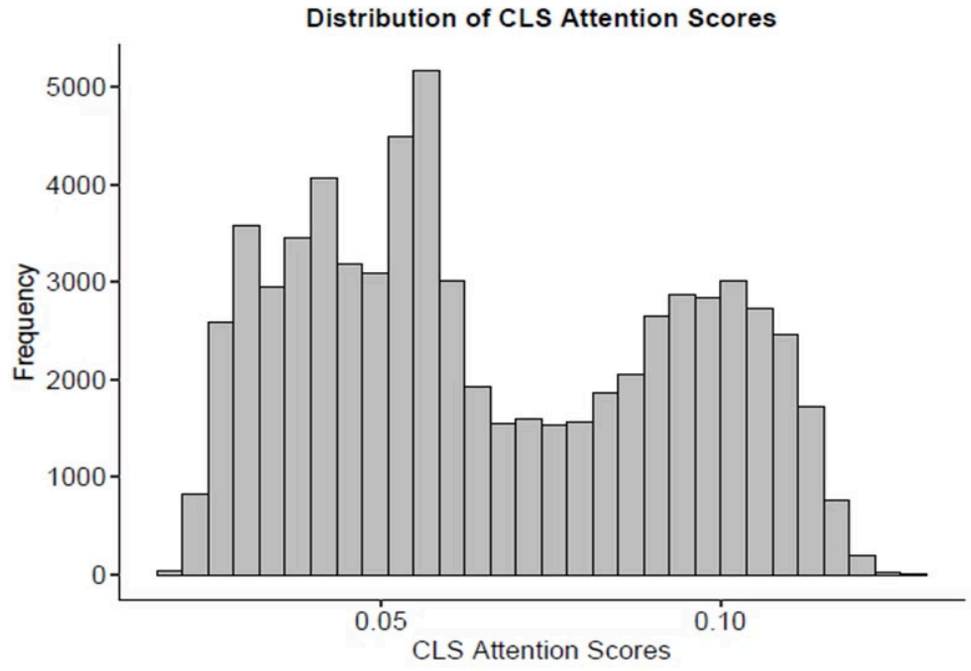

AI can be adapted to produce actuarially sound outputs

- Another example relates to mortality modelling and forecasting using deep neural networks
- Raw model outputs often are not suitable for direct use in actuarial work products
- Can we teach neural networks to conform to our preferences?
- Idea can be linked to Reinforcement Learning from Human Feedback, used to coax LLMs to produce acceptable outputs
- Deep Learning with Whittaker and Henderson shows how to incorporate smoothing into neural nets

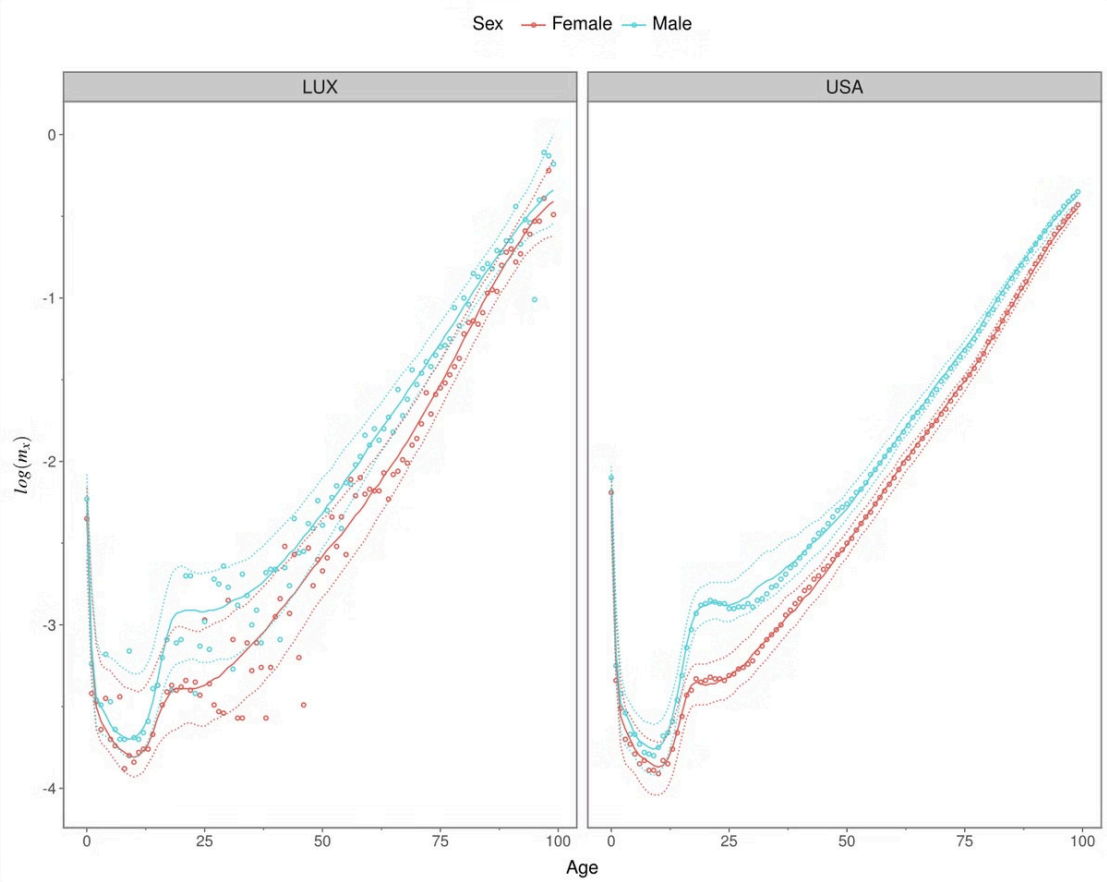

Model Explainability

1

Deep Learning Challenge

Deep learning models are often seen as "black boxes", making it difficult to understand how decisions are made.

2

Post-hoc Interpretability

Techniques like SHAP assign importance values to features, while PDPs and ICE plots visualize the relationship between features and outputs, aiding in model understanding.

3

Inherently Interpretable Models

Transformer models contain rich information about how they make decisions, allowing actuaries to understand and explain model behavior.

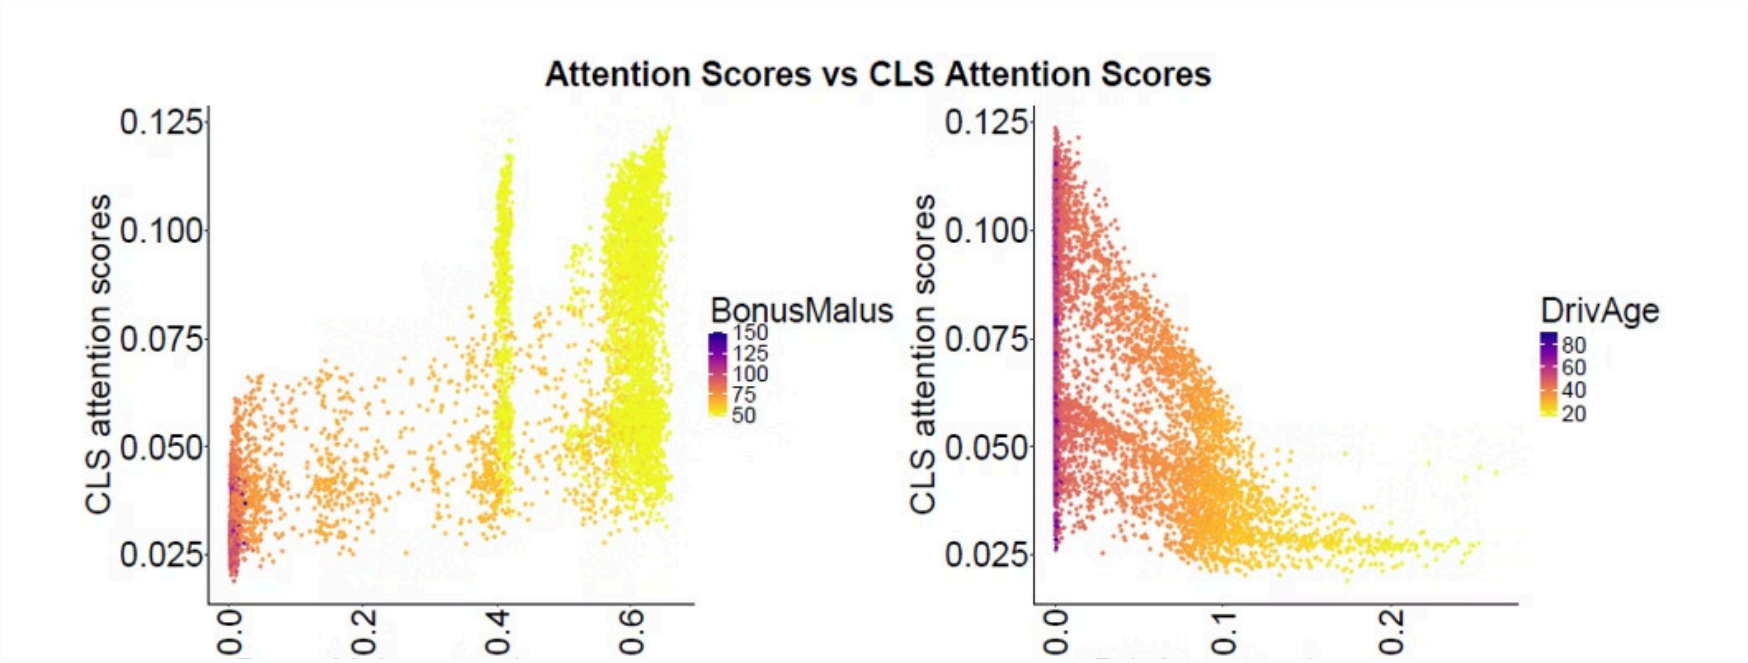

Mitigating Bias and Ensuring Fairness

1

AI Model Bias

AI models may inadvertently introduce bias, which can lead to discrimination against protected classes.

2

Discrimination-Free Pricing

Techniques like DFIP remove the impact of protected characteristics from insurance pricing models, ensuring fair pricing.

3

Multi-task Networks

Multi-task networks can produce hundreds of outputs at once, extending DFIP to more realistic scenarios, handling incomplete data while maintaining fairness.

Addressing Variability and Reproducibility

1 Neural Network Variability

Neural networks incorporate randomness, which can lead to variability in predictions across training runs.

2 Network Aggregation

Training multiple networks independently and averaging their predictions, known as nagging, reduces variability and enhances model reliability.

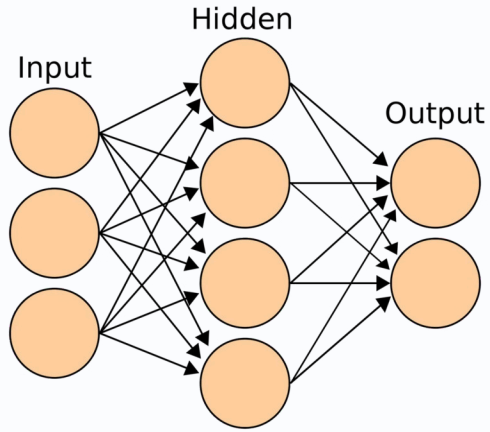

3 Reproducibility of Results

Nagging ensures the reproducibility of results, a critical aspect of deploying deep learning models in actuarial applications.

Key Takeaways

Adapting Deep Learning

Deep learning can be adapted to meet the specific needs and requirements of actuarial work, enhancing the profession's analytical capabilities.

Explainability Achieved

Model explainability can be achieved through post-hoc interpretability methods or inherently interpretable models, ensuring transparency.

Mitigating Bias

Bias and discrimination risks can be mitigated using techniques like discrimination-free insurance pricing and multi-task networks, promoting fairness.

Enhancing Reliability

Addressing variability and reproducibility issues through network aggregation ensures the reliability of AI models in actuarial practice.

Professional Requirements and AI

Professional Standards Overview

Actuaries must navigate a landscape of professional standards and guidance, ensuring compliance with requirements related to model understanding, fitness for purpose, and bias avoidance.

Documentation and Governance

The importance of thorough documentation, clear communication, and robust model governance cannot be overstated in the context of AI and deep learning.

Model Compliance

Ensuring AI models comply with professional standards is essential for maintaining the integrity and trustworthiness of actuarial work.

Implications for AI and Deep Learning Models

1 Compliance with Standards

Actuaries must ensure that their use of AI and deep learning models complies with relevant professional standards, focusing on model understanding, data quality, assumption justification, and model governance.

2 Documentation and Communication

Clear documentation and communication of model limitations and risks are paramount for ethical AI model use, avoiding bias and discrimination.

3 Collaboration for Best Practices

Staying up-to-date with emerging best practices and collaborating with other professionals is necessary to ensure compliance and the responsible use of AI in actuarial work.

Model Understanding

1

TAS 100 Principle 5 (Models)

Actuarial models must be fit for purpose, understandable, and have sufficient controls and testing.

2

The Actuaries' Code (UK)

Emphasizes competence and care, requiring relevant knowledge and skill.

3

APN 901 (South Africa)

Mandates understanding of models, including intended purposes and limitations.

- The TAS 100 Model Guidance explicitly requires understanding AI/ML models to comply with TAS 100 Principle 5.1.
- Challenges in understanding AI/ML models are acknowledged, especially for complex machine learning models.
- Alternative approaches for understanding complex models are suggested by the TAS 100 Model Guidance.
- Model validation, crucial for effective model risk management, is highlighted by actuarial standards and guidance.
- Regulatory bodies and professional organizations emphasize the need for model interpretability and transparency in actuarial work.

Key Takeaways

Established Standards

UK and SA have established professional standards and guidance governing actuarial work, with key principles focusing on professionalism, competence, integrity, compliance, and communication.

Model Guidance

Specific guidance on model understanding, validation, and governance is provided to navigate the complexities of AI and deep learning models.

Implications for Actuaries

Actuaries must consider the implications of these standards when using AI and deep learning models, ensuring that their work maintains the integrity and reputation of the profession.

Value Creation and the AI-Enhanced Actuary

1. ## Society Benefits

   AI-enhanced actuaries contribute to more stable and resilient insurance markets, benefiting society as a whole.

2. ## Policyholder Advantages

   Policyholders enjoy personalized products, fairer premiums, and an improved customer experience thanks to AI.

3. ## Insurer Value

   Insurers gain from increased efficiency, innovation, and data-driven decision-making enabled by AI.

4. ## Actuarial Opportunities

   Actuaries have the opportunity to enhance their skills, take on leadership roles, and drive innovation in the field.

Potential Future Roles for Actuaries in AI

Unique value that actuaries can bring to the development, governance, and application of AI systems **outside of the traditional actuarial domains**. By combining

- understanding of risk
- commitment to professional and ethical standards
- newly acquired AI skills

actuaries can play a pivotal role in shaping the future of AI and ensuring its responsible and beneficial deployment.

But only if we can deal with challenges in our own back yard.

Example Future Roles

1. AI Risk Assessment and Management Specialist

- Develop AI systems for risk assessment and management beyond the traditional insurance domain
- Apply actuarial expertise in risk quantification, coupled with AI skills, to fields such as finance, healthcare, and climate risk management
- Design and validate AI models that assess and mitigate risks in complex domains, ensuring accuracy, transparency, fairness, and alignment with regulatory requirements
- Contribute to the development of AI risk management frameworks and best practices

2. AI Governance and Ethics Consultant

- Contribute to the development of AI governance frameworks and ethical guidelines
- Bridge the gap between technical AI development and ethical considerations, ensuring responsible AI deployment
- Leverage actuarial expertise in balancing competing interests and maintaining professional standards
- Advise organizations on implementing AI governance structures, auditing AI systems for fairness and transparency, and managing AI-related risks
- Collaborate with policymakers, regulators, and other stakeholders to shape AI governance policies and regulations

3. Resilient AI Systems Architect

- Contribute to the development of AI systems that are resilient and adaptable to changing conditions
- Leverage actuarial expertise in stress testing, scenario analysis, and model validation
- Design AI systems that can handle uncertainty, adapt to evolving risks, and remain resilient under extreme stress
- Collaborate with AI engineers and domain experts to integrate resilience considerations into AI system architecture and deployment
- Develop frameworks for testing, monitoring, and maintaining the resilience of AI systems over time

Educational Considerations

1

Current State of Syllabi

The current state of AI and deep learning in actuarial syllabi is under review, with proposed changes to incorporate these concepts more thoroughly.

2

Transformation in Progress

Changes to the syllabus may include a new subject on machine and deep learning, tailored specifically for actuaries.

3

Resources Ready

Many of the resources needed for a new actuarial syllabus are freely available! See reading list for more details.

Example syllabus (see Harris et al.)

- Introduction to Machine Learning and its Actuarial Applications
- Supervised Learning for Actuarial Modelling
- Unsupervised Learning for Actuarial Analysis
- Introduction to Neural Networks and Deep Learning for Actuarial Applications
- Modifying Deep Learning for Actuarial Purposes
- Handling Large Datasets in Actuarial Work
- Ethics and Privacy in Actuarial Machine Learning
- Specific Applications

Positive and Negative Scenarios for the Actuarial Profession

The Thriving AI-Enhanced Actuary

- Rapid adaptation to AI and machine learning through modernized education
- Innovation in applying AI to traditional actuarial problems
- Data-driven critical thinking with ethical AI implementation
- Leading industry innovation and shaping the future of insurance

The Risk of Falling Behind

- Resistance to embracing AI and deep learning tools
- Loss of technical domains to data scientists
- Limited focus on regulatory matters only
- Diminished influence in the evolving insurance landscape

The Vision of the AI-Enhanced Actuary (1)

- Rapid advancements in AI and deep learning **could** revolutionize the actuarial profession.
- **How?** AI-powered actuarial models can provide more accurate insights, incorporate new data sources, and automate routine tasks.
- **AI-enhanced actuaries** can unlock new opportunities for innovation, efficiency, and value creation in traditional actuarial tasks.
- The vision of the AI-enhanced actuary is about harnessing AI to **amplify** the unique strengths and values of the actuarial profession.
- Challenges along the way include model explainability, regulatory hurdles, bias and discrimination risks, and the need for new skills and knowledge…
- … but **solutions already exist.**

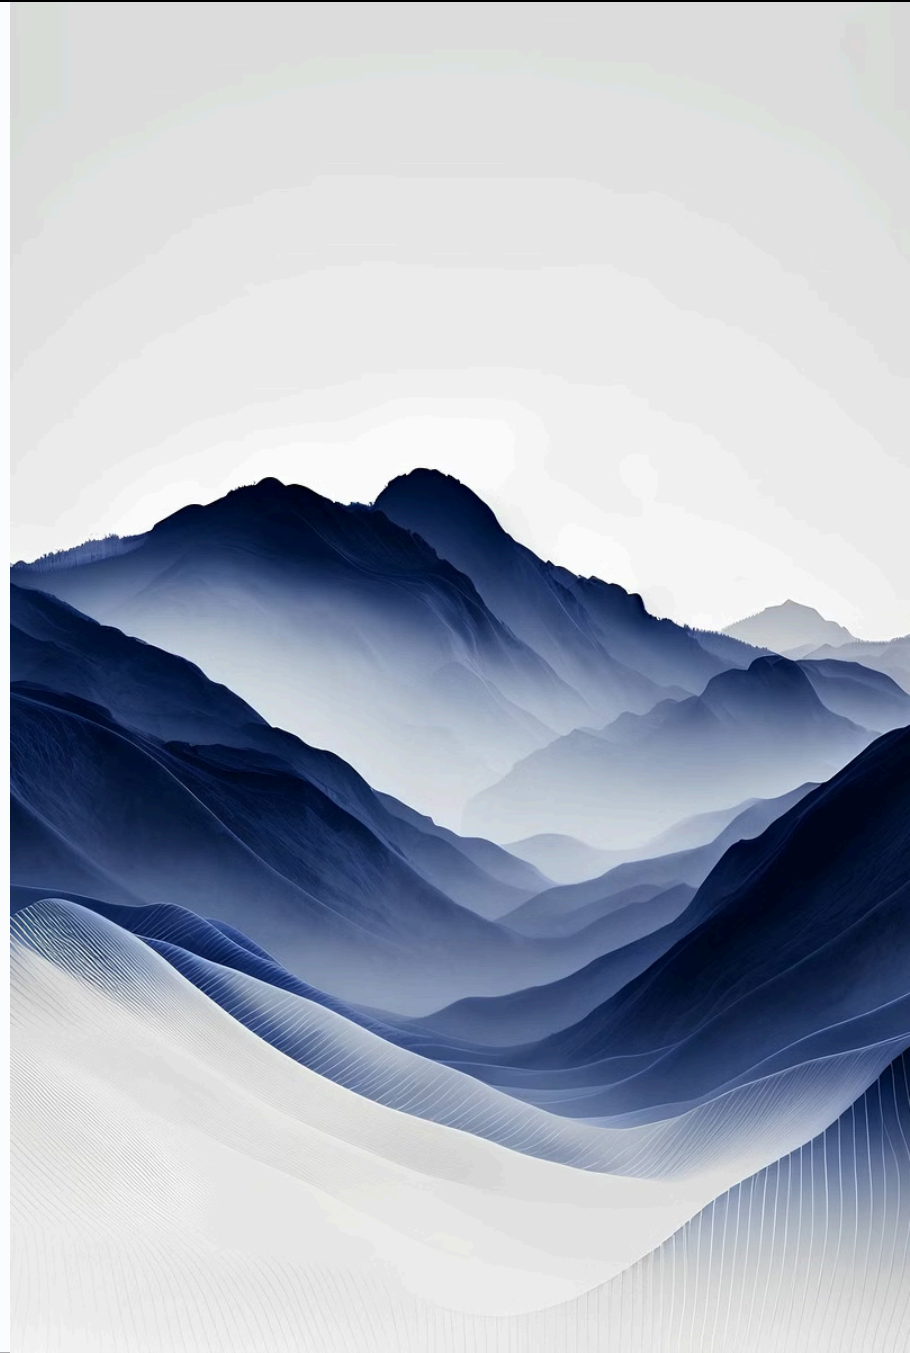

The Vision of the AI-Enhanced Actuary (2)

- Proactively addressing these issues through technical solutions, stakeholder engagement, ethical safeguards, and educational initiatives is crucial.
- Benefits include more stable insurance markets, personalized products, fairer premiums for policyholders, increased efficiency for insurers, and new career opportunities for actuaries.
- **Actuaries also have the opportunity** to shape the development and application of AI beyond insurance, contributing to responsible deployment in finance, healthcare, and risk management.
- The **actuarial profession must evolve** by updating educational curricula, fostering a culture of continuous learning, and collaborating with other disciplines.
- **The future of the actuarial profession is bright, and it starts with how we take actuarial adaptation forward from here!**

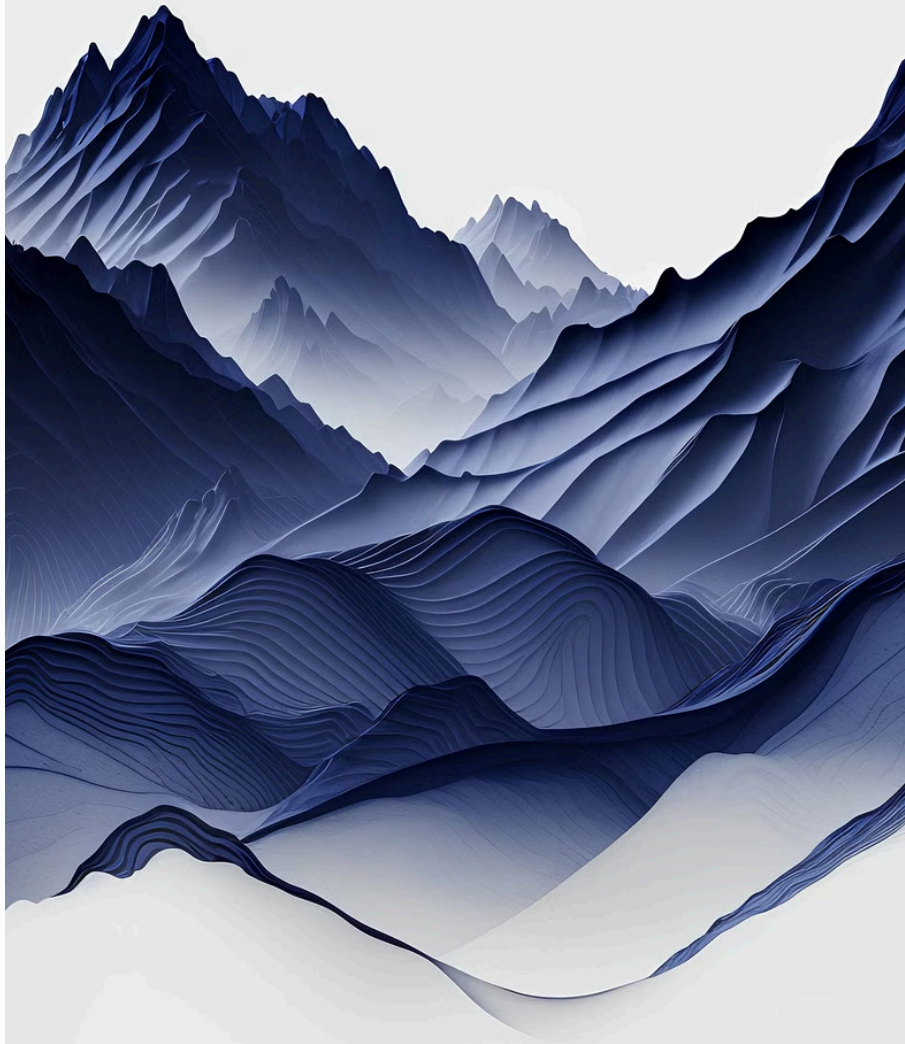

Reading List

- **Wüthrich, M. V., & Merz, M. (2023). Statistical foundations of actuarial learning and its applications. Springer:** A comprehensive book on the statistical foundations of machine learning and its applications in actuarial science, suitable for readers with a strong mathematical background.
- **Richman, R. (2021). AI in actuarial science–a review of recent advances–part 1. Annals of Actuarial Science, 15(2), 207-229:** A comprehensive review of AI and deep learning applications in actuarial science, providing a solid foundation for understanding the field.
- **Richman, R. (2021). AI in actuarial science–a review of recent advances–part 2. Annals of Actuarial Science, 15(2), 230-258:** The second part of the review, covering additional topics and developments in AI and deep learning for actuarial science.
- **Gabrielli, A., Richman, R., & Wüthrich, M. V. (2020). Neural network embedding of the over-dispersed Poisson reserving model. Scandinavian Actuarial Journal, 2020(1), 1-29:** A paper demonstrating the application of neural networks to reserving models, showcasing the potential of deep learning in this area.
- **Lindholm, M., Richman, R., Tsanakas, A., & Wüthrich, M. V. (2022). Discrimination-free insurance pricing. ASTIN Bulletin, 52(1), 55-89:** An important paper addressing the issue of bias and discrimination in insurance pricing models and proposing techniques for discrimination-free pricing.

Reading List (2)

- **Blier-Wong, C., Baillargeon, Jean-Thomas, C., Hélène, Lamontagne, L., & Marceau, E. (2021). Rethinking representations in P&C actuarial science with deep neural networks:** This paper discusses how deep neural networks can create representations useful for many different P&C insurance actuarial tasks.
- **Richman, R., & Wüthrich, M. V. (2023). LocalGLMnet: interpretable deep learning for tabular data. Scandinavian Actuarial Journal, 2023(1), 71-95:** A paper introducing the LocalGLMnet architecture, which combines the interpretability of GLMs with the predictive power of neural networks.
- **Richman, R. (2022). Mind the gap–safely incorporating deep learning models into the actuarial toolkit. British Actuarial Journal, 27, e21:** A paper discussing the challenges and considerations for incorporating deep learning models into actuarial practice while maintaining professional standards and ethics.
- **Schelldorfer, J., & Wüthrich, M. V. (2021). LocalGLMnet: A deep learning architecture for actuaries. SSRN:3900350:** Another paper on the LocalGLMnet architecture, providing further insights into its application in actuarial science.
- **Wüthrich, M. V. (2020). Bias regularization in neural network models for general insurance pricing. European Actuarial Journal, 10(1), 179-202:** A paper addressing the issue of bias in neural network models for insurance pricing and proposing regularization techniques to mitigate this problem.