

A Primer on Generative AI for Actuaries

February | 2024

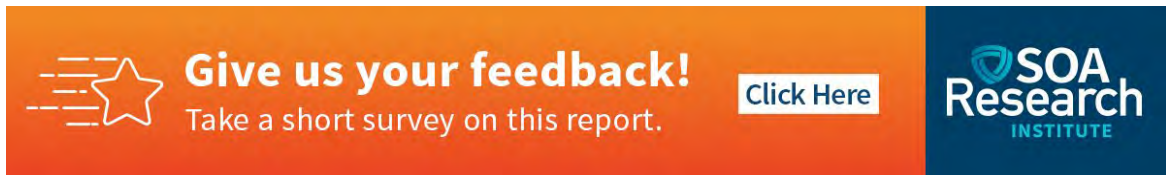




A Primer on Generative AI for Actuaries

AUTHORS Stephen Carlin, FIA
Stephan Mathys, FSA

SPONSOR Actuarial Innovation and Technology
Strategic Research Program Steering
Committee

A horizontal banner with an orange-to-red gradient background. On the left is a white star icon with horizontal lines extending from its left side. To the right of the star is the text "Give us your feedback!" in a bold, white, sans-serif font, followed by "Take a short survey on this report." in a smaller white font. Further right is a white rectangular button with the text "Click Here" in a blue, sans-serif font. On the far right of the banner is the SOA Research Institute logo, which consists of a blue shield icon and the text "SOA Research INSTITUTE" in white and blue.

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, 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 © 2024 by the Society of Actuaries Research Institute. All rights reserved.

CONTENTS

Executive Summary	5
Section 1: Introduction	6
Section 2: An Introduction to Generative AI	7
2.1 Overview of Generative AI	7
2.2 Potential Benefits of Generative AI for Actuaries	7
2.3 Examples of Where AI Can Bring Significant Benefits:	8
Section 3: Applications of Generative AI for Actuaries	10
3.1 Generative AI for General Productivity	11
3.1.1 Meetings	11
3.1.2 Summarizing Documents	11
3.1.3 Learning Tools	12
3.1.4 Drafting Content	12
3.2 Generative AI for Coding and Software Development	13
3.2.1 Copilot Tools	13
3.2.2 Code Conversion	13
3.2.3 Building Actuarial Models	14
3.3 Model Documentation and Governance	14
3.4 Enriching, Manipulating and Analyzing Data	16
3.4.1 Enriching Data	16
3.4.2 Manipulating Data	17
3.4.3 Data Analysis	18
3.5 Scenario Analysis	19
3.6 Automation and Efficiency	21
3.6.1 Automation	21
3.6.2 Efficiency	21
3.7 Claims	22
3.8 Underwriting	24
Section 4: Practical Application of Generative AI	26
4.1 Limitations and Areas of Concern When Using AI	26
4.2 Architectural and Deployment Considerations	27
4.2.1 Models and Platforms	28
4.2.2 Training Process and Data	30
4.2.3 Prompt Engineering, Transfer Learning and RAG	30
4.3 Checklist for Using Generative AI	33
Section 5: Summary (Conclusions)	35
Section 6: Acknowledgments	36
Appendix A: Summary of Key GenAI Models and Techniques	37
Neural Networks and Deep Learning	37
Transformer Models	37
Variational Autoencoders (VAEs)	37
Generative Adversarial Networks (GANs)	38
Appendix B: Glossary of Terms	40
Generative AI	40
Data Science	41
Machine Learning	41
Underlying Models and Techniques	41
Computing and Technology Terms	42
Broader Types of AI	42
Interacting with Generative AI Tools	42
References	44

About The Society of Actuaries Research Institute 46

Executive Summary

Generative Artificial Intelligence (Generative AI) builds on the core concepts of artificial intelligence (AI) to generate, or create new data or material. This generated output can be in the form of text, images, or even computer code.


To create a Generative AI model, vast amounts of data are used to “train” the system on the most likely response to a given prompt. The models then use either structured or unstructured learning processes to develop their predictive capabilities. An underlying Generative AI model will be exposed to end-users through an interface, where the user will prompt the model to produce an output. Users may have the flexibility to prompt the model in different ways, resulting in different types, volumes, or styles of output.

Actuarial applications of Generative AI may emerge in several areas. General task automation (eliminating manual tasks such as meeting note-taking) is an immediate increase to productivity. Generative AI may also be applied to many areas. General tasks include training employees, summarizing results, generating code, or understanding and interpreting model code (especially for older or less well-known code languages). Actuarial tasks include acting as a co-pilot to support an actuary in developing a tool or process, model documentation, creating audit trails for model governance, data enrichment (creating synthetic test data), data validation, automatic report analysis, scenario development, understanding the chain of data, peer review, claims management, and even underwriting.

Actuaries must also be cognizant of several limitations when applying Generative AI. These include challenges with quality and accuracy; repeatability; data privacy; auditability; copyright (for models trained on public or semi-public databases); training data quality; fraudulent use; the necessary skill of prompt engineering in order to derive useful outputs; and ethical considerations regarding bias, transparency, accessibility, accountability, and regulation.


Before deploying any Generative AI application into a production process, consideration should also be given to the architecture and accessibility, whether it will be in-house or a commercial external host. Additional considerations include the underlying model and training data required, the applicability of that training data to the specific task, and whether any feedback loop (where the model “learns” from the prompts supplied) will be incorporated.

This paper offers a technical overview of the mechanics of Generative AI, with several actual examples throughout. Some links to conversations with ChatGPT are included. It is not necessary to have a ChatGPT subscription to view these examples and, for the sake of brevity within this paper, they are not repeated within. This research paper is intended to be a practical introduction to the concept, not a comprehensive view. Actuaries should consider this a point-in-time view of this rapidly-evolving industry and use the issues presented here to spark conversation within organizations and across the profession.



Give us your feedback!
Take a short survey on this report.

[Click Here](#)



Section 1: Introduction

Generative Artificial Intelligence (Generative AI or GenAI) is a rapidly evolving field with a dramatic rate of change in both the underlying tools and methodologies, as well as commercial applications. This is a relatively light “primer” paper that could be updated or supplemented in the future to reflect changes in ‘state of the art.’

AI has great potential to reduce the cost of cognitive and creative actuarial work in several areas. This report suggests areas of potential benefit and encourages individual actuaries to explore the opportunities Generative AI offers to enhance individual work, company performance, and the actuarial profession.

The report is structured around three themes:

- An introduction to Generative AI
- Potential applications of Generative AI for actuaries. This section draws on both applications in other fields that may have parallels in actuarial work, and direct analysis and case studies
- Practical considerations for using Generative AI

We also provide a glossary of important terms and some further explanation of key topics in the appendices at the end of this paper.

There are many possibilities for how Generative AI may affect actuaries and the broader insurance industry. The Actuarial Association of Europe provides some different perspectives: <https://actuary.eu/wp-content/uploads/2024/01/What-should-an-actuary-know-about-Artificial-Intelligence.pdf>.

In addition, the UK Government Actuaries Department, alongside the Financial Reporting Council, conducted a survey of (mainly UK) actuaries and their use of AI in early 2023, which provides some additional thoughts: https://media.frc.org.uk/documents/Research_on_the_use_of_Artificial_Intelligence_and_Machine_Learning_in_UK_actuarial_work_AK5H1We.pdf

Section 2: An Introduction to Generative AI

2.1 OVERVIEW OF GENERATIVE AI

The concept of artificial intelligence (AI) has been around since the 1950s and has evolved rapidly since then, somewhat linked to the availability of computing power. There are a number of branches of AI focusing on different techniques and applications:

- Machine Learning (ML): Algorithms that learn from data to make predictions or decisions
- Natural Language Processing (NLP): Enabling computers to understand and process human languages
- Robotics: Integrating AI into physical machines
- Computer Vision: Enabling computers to interpret and understand the visual world
- Expert Systems: AI systems that emulate the decision-making ability of a human expert
- Neural Networks and Deep Learning: Advanced algorithms modeled after the human brain

Much of the recent wave of interest centers on Generative AI. Generative AI refers to algorithms that create (“generate”) new content, whether that be in the form of text, images, audio, or computer code. Unlike models, which simply categorize or label input data, generative models learn to generate new data. They capture the underlying patterns, structures, and distributions in the training data, enabling them to produce novel content that is similar to the original dataset. Generative AI is a broad category encompassing many different techniques, including large language models, deep fakes, and text-to-image models.

Large language models (LLMs) are a common and specific type of generative model designed to produce human-like text. They are trained on vast amounts of text data and can perform tasks like translation, summarization, question-answering, text generation, and a range of tasks relating to software coding. They are a subset of Generative AI, specifically focusing on natural language processing (NLP). While LLMs are only one subset of Generative AI, they are relevant to many of the applications we discuss and give particular attention to in this paper.

Transformer models are a common architecture for LLMs - particularly Generative Pre-trained Transformer (GPTs). Transformer models are trained to predict the next word in a sentence given all the preceding words. Chat GPT - developed by OpenAI - is one of the best known Generative AI platforms and uses this approach, though there are now many models and platforms available from a variety of providers.

The capabilities of generative models have significantly advanced in recent years. Such advances have been driven by improvements in the available algorithms and also the substantial expansion of available computing power through cloud computing platforms.

2.2 POTENTIAL BENEFITS OF GENERATIVE AI FOR ACTUARIES

There are certain tasks Generative AI tools excel at - particularly tasks that require processing vast amounts of data, identifying patterns, and generating content at a scale and speed challenging for humans. For actuaries, Generative AI could be leveraged to unlock new capabilities or deliver significant efficiency gains.

2.3 EXAMPLES OF WHERE AI CAN BRING SIGNIFICANT BENEFITS:

- **Coding-related Tasks:** Generative AI tools can be trained on vast repositories of code, allowing them to perform a number of tasks relating to coding and software engineering, such as generating code based on natural language prompts, reviewing and improving code, generating test cases and code comments. For actuaries with advanced coding skills, these tools offer potential productivity gains. Perhaps just as importantly, they offer a boost to actuaries with more limited coding skills by accelerating the learning process.
- **Data Processing and Analysis:** AI can analyze and process data on a substantially larger scale than typical humans, and in ways which draw new insights with a range of applications for actuaries and insurers.
 - **Pattern Recognition in Large Datasets:** AI algorithms are adept at identifying complex patterns and correlations in data that might be too subtle or complex for human detection. This has applications in data validation and fraud detection.
 - **Predictive Modeling:** AI models can forecast trends and outcomes based on historical data, drawing out additional insights or more accurate forecasts than some traditional methods. For actuaries, this has a potential application in policyholder behavior prediction or new forms of experience analysis.
- **Content Summarization and Generation:** AI can process large volumes of input data and generate large volumes of output content, such as text or images, in a fraction of the time it would take a human. Application for actuaries includes generating documents and reports or summarizing large volumes of text, such as current or proposed regulations. This may also lead to finely-tuned chatbots, which can respond to policyholder queries with appropriate content.
- **Simulations and Scenario Generation:** AI can generate a multitude of scenarios or simulations rapidly, which is invaluable in actuarial modeling applications like climate modeling, stress testing, and scenario modeling.
- **Consistency in Repetitive Tasks:** Unlike humans, AI does not tire and can perform repetitive tasks with consistent accuracy, which is crucial in quality-controlled processes, with potential applications for automating recurring tasks in actuarial workflows.
- **Image Generation and Processing:** AI can generate new realistic images based on input prompts. It can also be used to process images and have been used in claims processing to assess damage in photographs in home or auto insurance.
- **Customization and Personalization:** AI can analyze individual preferences and behaviors to provide personalized recommendations, content, or services - something impractical for humans to do on a large scale. AI has the potential to enable new types of insurance products with customizable features and pricing that can be offered to specific customers.

Generative AI tools offer powerful capabilities in a number of areas directly relevant to actuarial work, but it's important to note they require human guidance to set objectives, interpret results, and provide oversight. The combination of AI's computational power with an actuary's domain knowledge, creativity, contextual understanding, and ethical judgment may yield results substantially more valuable than either acting independently.

There are several areas of concern and limitations associated with the use of Generative AI. We explore these in detail later in the report. For now, we note three areas of concern:

- **Actuarial judgment** – While AI algorithms can appear to show remarkable skills in comprehension, they are fundamentally probabilistic models. As such, they are no substitute for actuarial expertise

and judgment and should be viewed as an augmentation – not a replacement - for the entire portfolio of actuarial skills, which actuaries apply to specific concerns.

- Accuracy - Generative AI tools, particularly LLMs, can occasionally produce plausible-sounding but misleading or erroneous results. This should be factored into the way Generative AI is used for important tasks.
- Security - Users must be aware of how Generative AI tools treat data from a security and privacy standpoint and validate whether the platform, model, and application comply with internal data policies and legal requirements.

We explore these, and other concerns, later, along with strategies for mitigating any limitations.

Section 3: Applications of Generative AI for Actuaries

Applications of Generative AI are broad and growing quickly - hardly a day goes by without a new application of AI in the news. Historically, much of the focus for AI has been on either:

- Supplementary modeling in science and technology fields - e.g., using AI to simulate outcomes rather than running expensive tests, or simulating new molecular structures for drugs; or
- Replacing manually intensive but 'low skill' jobs (e.g., replacing call centers with AI-powered chatbots).

One of the major shifts recently has been the focus on roles and industries previously considered 'immune' to AI, like creative roles (marketing), technology (programming), and other white-collar industries.

In creative fields, the generation of new or manipulated images and videos - once the preserve of artists and Hollywood special effects teams - are available readily through personal devices and often for negligible costs. For example, AI services like Dall-E and Midjourney are able to create new images based on text-based prompts. Large language models are able to create articles and marketing copy with remarkable coherence and with customizable tone and personalization.

Similarly, the legal profession is going through a phase of disruption with the emergence of the LegalTech field, much of it powered by AI. AI can generate legal documents, such as contracts or agreements, by understanding the requirements of a particular case and using existing templates. This speeds up the drafting process and reduces errors. AI tools can analyze large volumes of documents for due diligence or discovery processes. They can identify relevant passages, inconsistencies, or potential issues much faster than manual reviews. AI can assist in legal research by quickly sorting through vast legal databases, identifying relevant case laws, statutes, and legal precedents, thus saving considerable time for legal professionals.

Even just a few years ago (see, for example, McKinsey 2017), the expectation was the adoption of AI would likely create demand for additional jobs in technology and software development. However, the effectiveness of GenAI in creating, reviewing, documenting, and testing software code potentially offers a step change in development productivity. We discuss this further in subsection 4.2.

AI is being widely deployed in fraud detection. Generative AI models are capable of analyzing transaction patterns to identify unusual activities that may indicate fraud. These systems can learn from historical data to recognize signs of fraudulent transactions in real-time, as well as identify indicators of activities like money laundering, identity theft, or account takeovers.

In insurance, Generative AI can be used across claims processing, underwriting and customer service:

- AI models can automate much of the claims processing work. This includes traditionally time-consuming tasks like data entry, document scanning, and initial assessments.
- Fraud and anomaly detection: AI systems are trained to identify patterns commonly associated with fraudulent claims. This includes anomalies in billing, unusual claim frequencies, or inconsistencies in claim details.
- In auto or property insurance, AI algorithms can analyze photos or videos of damage to provide initial estimates, speeding up the claim approval process.
- Instead of relying on broad categories, AI enables more personalized underwriting. Policies can be tailored to individual risk profiles, considering factors like exercise habits, dietary patterns, and even genetic information.

- AI can generate automated responses to customer inquiries, providing instant assistance and guidance throughout the claims process.

Generative AI offers huge potential to boost productivity and bring new capabilities to actuaries. In this section, we explore some potential applications in the actuarial domain and in the insurance industry generally. While these examples cover a number of themes, this list is not exhaustive. The reader may take inspiration from the examples below and seek out new opportunities.

3.1 GENERATIVE AI FOR GENERAL PRODUCTIVITY

Actuaries, like any other office-based professionals, typically spend a significant part of their day dealing with emails, attending meetings, and writing reports and memos. Generative AI increasingly offers solutions to streamline some of the less analytical parts of the job to free up time to focus more on the value-added tasks, such as analysis and consideration of implications.

3.1.1 MEETINGS

AI powered transcription tools are able to record meetings with remarkable accuracy - including the ability to differentiate between different speakers. These tools typically also offer the ability to create meeting summaries that distill key discussion points and actions. For those of us who are unable to type fast enough to keep up with a conversation and rely on typing up hastily scribbled notes after the event, this can be a huge timesaver (as well as preserving details that may be lost in short-form notes). A range of dedicated note-taking platforms (such as Fireflies.ai) exist, but this functionality is increasingly embedded into virtual meeting applications such as Microsoft Teams and Zoom.

It is clear that concerns about data privacy and security apply, as well as any issues around the consent required for recording. There may also be a specific concern around recording meetings in such detail - sometimes we may not want everything we say recorded for posterity. This comes with its own risk: that it discourages full and frank discussions, which may be beneficial to the situation. However, overall, this can be a significant gain through saved time and more effective meetings.

3.1.2 SUMMARIZING DOCUMENTS

Large Language Models can be effective in summarizing large documents or other bodies of text. This typically requires the uploading of one or more documents directly or via an API - features not available with the free tier of many GenAI platforms. However, once uploaded, the LMMs can give great summaries of the overall themes in a document, or answer specific questions about the content.

These summaries can be very sensitive to the prompts used. While no substitute for a thorough reading of the document, they can be useful in an initial scan of literature where the comparison is not a cover-to-cover read, but a quick scan of the contents and a skim of the executive summary. They can also be helpful in addition to a thorough read, helping recall specific sections or collating related themes through disparate sections of a note.

They can also be very useful in identifying and summarizing differences between two versions of a document - this can be helpful when analyzing drafts of lengthy documents, such as different drafts of rules and regulations.

Again - use of AI in this context is no substitute for a detailed read of necessary and relevant documents - but quickly processing large volumes of text is a key capability of Generative AI tools that actuaries can use to their advantage when appropriate.

3.1.3 LEARNING TOOLS

Generative AI can be used in a number of different ways to facilitate learning. It is increasingly integrated into training and learning platforms to provide customized content, interactive questioning and response, as well as the ability to provide automated feedback on quizzes and assignments.

Over the past 20 years or so, Google and Wikipedia have been the go-to resources for people seeking information about a subject. With their vast sets of training data, Generative AI tools offer an alternative source - a key advantage being the ability to craft an answer specific to the question asked and the ability to ask follow-up questions for clarity or additional detail. This can help speed up the process of obtaining a basic understanding of a topic.

There are some caveats. The risk of GenAI “hallucinations” (where the GenAI model invents untrue facts and reference material) means one should never take the generated information as gospel - particularly for important tasks. Importantly, and as with any model, the output is only as good as the input. At the time of writing, pretrained Generative AI models often have significant data lags - i.e., they are trained on data up to a specific point in time prior to the publication of the model. This limits the usefulness of the approach for recent news or events.

Although the datasets are large and diverse, any given model is likely to contain limited information about very specific technical topics (such as those encountered in actuarial work), which may limit the value for such tasks. However, actuaries have encountered numerous examples where the outputs were surprisingly helpful for what were expected to be relatively obscure actuarial topics.

GenAI can also be an extremely powerful tool in helping improve coding and software development skills - this specific topic is addressed later in subsection 4.2.

3.1.4 DRAFTING CONTENT

A key and defining capability of Generative AI - particularly LLMs - is the generation of new text-based content such as emails, reports, marketing collateral, and presentations.

The quality, depth, and appropriateness of the content will be dependent on the quality of the prompts and the relevance of data included in the training or augmentation datasets. However, even if the GenAI outputs only provide some useful concepts and structure users can adapt and extend - this can be a huge labor saver, particularly for those who struggle when faced with a blank sheet of paper or empty presentation.

For example, the authors made significant use of Generative AI (mainly GPT-4) to assist in the creation of some parts of this report - particularly for ideation, rough drafting, and background knowledge. The glossary of terms was largely produced by GPT-4, with some refinements and small manual additions. However, despite significant efforts, we were unable to come up with a single prompt that could create the entire document for us to a sufficient standard.

Another example of this is the use of Generative AI to create model documentation. We expand on this in a later section (subsection 4.3).

3.2 GENERATIVE AI FOR CODING AND SOFTWARE DEVELOPMENT

A key strength of Generative AI is its ability to interpret and write software code.

Actuaries' expertise with coding and development varies significantly with the nature of the role, but most actuaries fall into one of three categories:

- Actuaries who never code, such as management positions or commercial roles;
- Actuaries who are actively coding on a regular basis for model development, data analysis and manipulation, or process automation; and
- Actuaries who don't code - but would benefit from doing so to reduce the amount of Excel-based or manual manipulation of data and/or repetitive tasks.

Generative AI offers potentially significant benefits to actuaries in categories 2 and 3.

3.2.1 COPILOT TOOLS

The software engineering industry, generally at the forefront of embracing technology to make their lives easier, has been an early adopter of leveraging Generative AI. Increasingly, software development tools offer a range of integrated AI tools or plugins to accelerate the development process. Some of the key players are GitHub Copilot, Amazon Code Whisperer, and GPT-4. Even at this relatively early stage in their evolution, these tools boast productivity improvements of [50%-60%](#) with an expectation of significantly more to come.

These tools offer a range of capabilities to assist with the development process in a number of ways, each of which can accelerate the development process:

- Code generation based on text prompts;
- Code review, bug fixing, and suggested improvements;
- Generating code comments to explain and document code; and
- Generating unit and regression tests that can be integrated into the build process.

The first two can also be very helpful as learning tools - particularly when learning a new language. Even if the suggested code doesn't meet the exact requirements, it generally gives a reasonable starting point, like recording an Excel macro while performing specific tasks and then adapting that code to more general cases.

Similarly, code review suggestions can help users improve their coding skills by offering more elegant structures or refined syntax, not only improving the current code, but helping users develop their understanding of a language for next time.

The addition of comments to code, while a generally accepted best practice, is often neglected (or a token effort made). AI coding tools show a good ability to interpret code and provide commentary as to its purpose. While useful for code-level documentation, this can be extended to full model documentation, which we discuss in the next section.

3.2.2 CODE CONVERSION

One of the code-related tasks that Generative AI seems to be particularly strong at is converting code from one language to another - e.g., Python to Google Golang or vice versa.

For insurers, this has significant potential to ease the pain of migrating away from legacy systems. Many insurers have muddled through on very old systems because it has been easier and/or cheaper to patch-up an old system than face the challenges of migrating. This is particularly the case for old mainframe-based admin systems where code migration is only part of the challenge. However, there is an increasingly acute skills shortage in some of these languages (like APL) as the original generation of developers leaves the workforce.

While Generative AI tools do a very good job of converting syntax, challenges remain that require human input:

- Code efficiency - does the code need to be re-architected or refactored? Many of the constraints these solutions were originally designed for (limited RAM, single-threaded processing) no longer exist. Similarly, there may be the opportunity for code to be cleaned up to remove redundant variables and dead code.
- Proprietary code and copyrights - can the code for legacy systems be accessed at all? If it can be accessed, are there copyright issues preventing direct conversion?
- Proprietary syntax and functions - some systems also use proprietary languages that include proprietary functions. On the face of it, this should be a significant barrier. Experience with a GenAI model reading a proprietary scripting language (based on Google Golang) suggests that GenAI does a very good job of interpreting proprietary code, even without any details on the specific rules of that code language. With appropriate additional data used to augment prompts, the barrier to understanding proprietary functions can be eliminated.

3.2.3 BUILDING ACTUARIAL MODELS

What does the opportunity for automated code generation or conversion mean for actuarial modelers? Can one ask a Generative AI tool to, for example, write the code for a term life insurance reserving model and get a functional model written in the language of our choosing? Or can a Generative AI tool scroll through a library of code in a specific language and replicate the functionality in a different language? Not yet - but perhaps this type of first draft is not too far off.

[In this example](#), we ask GPT-4 to create a term life reserving mode. With no augmenting data, it is able to produce a reasonable, yet heavily-simplified, model. Such a simple model is typical of the examples often provided in examination study materials. This is likely an indication of the type and complexity of the texts included in the data used to train the model.

The use of Generative AI to build and reconcile actuarial models is an ongoing research theme across insurers and technology vendors. There is great potential to quickly build baseline models for further refinement, eliminating some of the tedious first steps. However, this requires rigorous engineering of prompts and curation of appropriate augmentation data and resources. Additionally, using Generative AI for initial model creation does not eliminate the need for appropriate testing and governance mechanisms once such code is deemed appropriate and deployed into production.

3.3 MODEL DOCUMENTATION AND GOVERNANCE

Anecdotally, actuarial models are generally not as well documented as audiences would like. Even when internal or regulatory initiatives drive efforts to create better documentation, this often falls into disrepair over time as maintenance of the documents may get overlooked or deprioritized in favor of business-as-usual production or projects with higher implied value.

Generative AI offers a potential solution to this - or at least a mechanism to lighten the load. The AI can be instructed to create a natural-language summary of functional code - essentially, translating from computer language to human language. This is closely aligned with the AI's ability to interpret code, as described above.

A simple demonstration of this is to give some model code to one of the tools and ask it to provide technical documentation for the model ([example](#) based on a term life model using GPT-4). With a bit of prompt engineering, it's possible to get a reasonably good document without needing additional augmenting data. In this case, the level of detail in the output was quite sensitive to the prompt used. With additional prompting, the AI will also include commentary on the appropriateness of the model, as well as a list of limitations.

The standard web interfaces for the AI tools might work for an experiment like the above, but deriving value in the long run requires a more structured approach. This is likely to involve connecting a specific prompting language and Retrieval Augmented Generation (RAG) process via an application programming interface (API). [RAG is the process to supply a repository of additional relevant information at the point of query, like supplemental information, which the AI can access if instructed to do so, and is discussed further in subsection 5.2.3.] This more structured approach is required to provide the specificity to ensure the documentation has the right level of detail and structure, but also to ensure consistency in the outputs in each execution. With the right investment in set-up, output can be made sufficient for real world purposes.

Practical example:

An actuarial software platform uses GenAI to create and maintain the documentation of the functions library supporting the model-scripting language. The company uses AWS Bedrock to regularly scan the functions library with instructions to interpret the functions found therein. The AI then refreshes the output documentation with any new or updated functions. It also creates a summary of changes, allowing the reviewing analyst to focus on the areas of change for the sign-off process. While the task is relatively simple - mainly focused on the inputs and outputs of the functions - the approach allows the documentation updates to be embedded in the company's continuous integration / continuous delivery (CI/CD) build process, meaning the documentation of functionality is always current.

Another recent example is the creation of real-time model documentation within a model development environment. [Such a model is built on the library of functions indicated above, using those functions to create specific calculations.] This uses a similar back-end approach (pretrained large language model and RAG model accessed via API) to generate model documentation as the model is developed or changes. Again, a key advantage of this kind of set-up is that documentation is automatically updated as the model changes - not a separate task that has to be separately scheduled with the risk that the model code and associated documentation diverge.

As well as 'snapshot' technical documentation, another major documentation activity is the creation of change and governance documentation. Again, with the right prompts and infrastructure, GenAI can do a useful job comparing and documenting the changes between model versions in a way that can be automated into the model development lifecycle. While not a substitute for human involvement, it is a very useful addition to the toolkit that excels in distilling complex changes in multiple locations into a digestible summary in a way that the Mark 1 Eyeball or even diff-compare applications don't.

A stylized example of this is shown in [this example](#). To illustrate the point, we used a term life model and a very similar model extended with a Total Permanent Disability (TPD) rider. Even with a basic prompt, GPT-4

is able to summarize the key difference between the two models (interestingly, it even infers the definition of TPD without any augmentation data).

Documentation is an area where Generative AI can have immediate benefits with relatively low barriers to implementation where access to (the model or other) code is available. This could be particularly useful where the code in question is written in a language where the skills may be in short supply - at least within actuarial - as is sometimes the case with older systems.

While our focus here has been on using GenAI to document code, the same philosophies (with different implementation) could be applied to data and assumptions. We explore some of this in the next section.

3.4 ENRICHING, MANIPULATING AND ANALYZING DATA

The application of Generative AI to data handling is a very broad field with a substantial volume of literature and research. For example, with data analysis, there is substantial overlap between Generative AI and more general AI applications, such as machine learning and predictive analytics. It is also an area with a rapidly increasing number of AI-powered tools designed to make these tasks easier, and/or perform back-end infrastructure work linked to the core GenAI systems. While we describe a number of themes below, these should be taken as illustrative rather than comprehensive.

3.4.1 ENRICHING DATA

Data enrichment is creating new data that looks like the training data. This is a key capability of Generative AI models. This may be necessary for testing calculation capacity or ensuring that the scope of all possible combinations of demographics, policy coverages, or policyholder options can be adequately processed by the system.

At the simpler end of the spectrum, this is something we have used extensively for generating test data when building and validating models. It's fairly easy to provide the AI tool with a .csv file containing model points or scenarios and ask it to create additional rows of data. For developing massive volumes of test data, Generative AI does a good enough job. In [this example](#), we gave GPT-4 a model point file with 10,000 rows of data and asked it to create a new version with 100,000 rows. Compared to a manual copy-paste exercise in a spreadsheet, this took less time and is easily scalable to substantially larger volumes.

GenAI can also be used to create synthetic datasets that reproduce the characteristics of customer data. This can be helpful where privacy concerns about using actual policyholder data restrict its usage.

Perhaps a more interesting use case, given the sparse nature of some insurance datasets like claims, is directing Generative AI to synthesize more data to support the application of a broader array of analytical techniques. Such techniques may even include machine learning models. These techniques are widely used in other fields for a variety of applications. Whether it is appropriate for similar applications in insurance will depend on the context. If used to generate additional test data for ML algorithms for processing claims or underwriting, synthetic test cases can be extremely valuable. However, actuaries should not generate synthetic claims data for experience analysis. Including synthetic claims data would imply greater statistical reliability than actually exists.

Another potentially significant application is the generation of scenario data for model and stress testing. We explore this in detail in the following subsection (4.5).

Overall, this is a field with potentially significant implications in insurance. While there are some quick wins in the easy things, much work is required to unlock the potential in more advanced applications.

3.4.2 MANIPULATING DATA

Most actuaries, at some point in their careers, have spent a significant amount of time manipulating data and getting it in a usable format. As a profession, it is something we spend a significant portion of our time on. Given our reliance on data, it's not unreasonable that we spend time getting it in the right shape, yet significant variation in our individual effectiveness at such preparation remains.

Already, a wide range of tools exists that can be used for this task, like R, Python and SQL. Several no-code / low-code data management solutions like Alteryx and AWS Glue are available. Many actuarial systems have inbuilt data handling modules, too.

Additionally, pre-installed spreadsheet software (Excel) and browser-available spreadsheet applications (Google Sheets) are easily available and, therefore, often used. However, such blank-slate applications require substantial development and often do not provide the same kind of capacity as purpose-built applications.

How can Generative AI help with data or file manipulation? AI-powered tools can do a reasonable job of converting data files to a required format based on a sufficiently descriptive prompt.

Example using Julius.ai

Economic Scenario Files are often formatted with the variables in columns, while the scenario identifiers and timesteps are stacked down the rows (see below). While common, it's not an easy format to analyze. Can Generative AI create more "usable" files?

Given the prompt below:

This file contains stochastic economic scenarios. The first and second columns contain time indicators (year and month). The third column contains the scenario number. Columns 4 through 16 contain values of economic variables for each timepoint for each scenario.

Can you split the file so that each economic variable is in its own file? Can you reformat the data such that the time dimension is across the columns with the scenarios in rows?

The Julius application did, indeed, return downloadable files with the correct formatting within a couple of minutes. It also created data validations and produced summary statistics. This may seem like no big deal for data-savvy actuaries with access to the right tools, but for those without access to the requisite skills and/or tooling, GenAI can be a catalyst, saving time during one-off instances of data manipulation.

For recurring usage, some degree of investment is required to develop tooling or processes to ensure outputs are produced consistently and efficiently. While this investment could be used to productionize the use of an AI tool, arguably the same effort could deliver workable solutions in Python or SQL. Such a solution may produce the same output without corresponding risks or concerns inherent in using AI-based tools. However, the AI may have been used to create the prototype Python code used to perform the manipulation, so even if AI is not used as the ongoing tool, it can provide a starting point for further work.

AI-powered tools clearly have some value as a utility for data-related tasks that might otherwise seem tricky. For actuaries with deep data and coding skills, the benefits might be less obvious in the short term, but there may still be value as a productivity tool and the rate of evolution suggests there might be much more to come.

3.4.3 DATA ANALYSIS

The role of AI in data analysis is another broad topic in its own right. It expands out beyond Generative AI to include a wide range of machine learning modeling, as well as data science more generally. This encompasses topics such as:

- Data validation
- Anomaly detection (e.g., used for fraud detection)
- Automated analysis and report production
- Predictive analytics and machine learning

We touch on these topics below. Detailed exploration of these themes is beyond the scope of Generative AI, but they are certainly areas of interest to actuaries and worthy of further research.

Data Validation and Anomaly Detection

Generative AI algorithms, like Variational Autoencoders (VAEs), are particularly adept at spotting issues and outliers in data. These algorithms can be trained to spot issues in different kinds of data, whether numerical, text, or mixed. They can also be trained to spot data entry errors, such as a policy inception date prior to a particular product's availability for sale.

VAEs are particularly effective for data validation due to their unique structure and the way they process and generate data. VAEs encode input data into a latent (hidden) space, representing data points as distributions rather than fixed points. This probabilistic approach allows for capturing the underlying uncertainties and variations in the data, making it ideal for validating real-world data where variations are common. By learning these distributions, VAEs can generate new data points similar to the training data. This ability is useful in data validation to create benchmark examples against which to compare new data. VAEs are trained to reconstruct input data from the latent space. Data points that are significantly different from the training data will have higher reconstruction errors. This property is useful in identifying anomalies or errors in new datasets.

VAEs are capable of learning complex, high-dimensional data patterns, which is essential for validating intricate datasets like those in finance or healthcare. Another useful characteristic is that VAEs can handle incomplete data by learning to infer missing values from the known ones, which is particularly beneficial in real-world datasets where missing data is common. While VAEs are often used for such work, other algorithms can achieve similar outcomes.

This ability to spot outliers can be a powerful tool in fraud detection - whether in banking transaction data or in insurance claims processing. Algorithms can be used to automatically flag transactions that show the indicators (or combinations of indicators) of suspicious behavior, such as the time of day a claim was submitted or the number of times the policy documentation was opened.

Automated Analysis and Report Production

The increasing proliferation of AI-powered data analysis tools has been much talked about for its potential to unlock insights previously unavailable to the average person on the street or business owner. There is certainly some merit in this claim, as the tools essentially lower the barriers of entry by eliminating the need to code or even have a proficiency in Excel.

For actuaries, the ability to extract basic statistics from a dataset doesn't exactly move the needle. In fact, in our early experiments with some of these tools, the speed and inefficiency of the English language in

expressing mathematical concepts was somewhat frustrating. However, there has been marked improvement in the capabilities of these tools over just a few months. It's likely that such tools will evolve and ultimately become part of the toolkit that forms part of the analysis tool-chain.

In the short-term, the benefit of Generative AI may be in lowering the barriers to entry to advanced analytics and machine-learning techniques.

Predictive Analytics and Machine Learning

This is an area of huge relevance to actuarial work that brings together Generative AI and the broader field of (not generative) machine learning algorithms. Before the explosion of interest in Generative AI, advanced analytics, data science, and machine learning had been the hot topics of the previous few years. While somewhat lost in the noise of late, this is an area of huge potential value that has not yet been embedded in the insurance industry.

These techniques can be used to power-up traditional analysis by including more data, distilling more powerful insights into the drivers of behaviors, or creating predictive models that can be used to proactively influence outcomes. Examples include:

- Advanced experience analysis - enriching experience analysis with additional socio-economic, health, or financial data, allowing the insurer to build up a more granular understanding of policyholders and what drives differences in experience across granular policyholder groups;
- Predictive models for lapse or claims that can support proactive intervention initiatives;
- Creation of models which predict the relationship between insurance prices and volumes sold;
- Analytics to support distribution quality management; and
- Model acceleration techniques such as the creation of proxy models that learn the behavior of cash-flow projection for valuation models, or advanced clustering algorithms for model point and scenario compression.

Again, detailed explorations of such applications are beyond the scope of this paper's concentration on Generative AI, but worthy of further exploration.

3.5 SCENARIO ANALYSIS

Scenario analysis, "what-if" analysis, or scenario planning can be enhanced by Generative AI to go beyond typical sensitivity analyses created by actuaries. For example, financial models may offer a best estimate of an insurance company's future financial performance. This forecast relies on hundreds of actuarial assumptions, as well as substantial assumptions about the external economic environment in which the company operates. These assumptions may include components such as investment markets, agent recruiting, regulations both already in-force and soon to be implemented, current client portfolio experience, and even demographic changes within the target market.

Typically, the best estimate forecast will be supplemented with a generic evaluation of the financial impact of one or several components shifted up or down from their baseline value. This might be Mortality +/- 10%, Lapses +/- 15%, and a series of defined interest rate paths. Such an analysis allows for the understanding of the sensitivity of company financials to each individual element. However, it does not offer any kind of evaluation of the impacts if more than one element deviates from the baseline at the same time. Nor does it help understand any impacts from nonlinearities in those elements. And, perhaps most importantly, any scenario analysis of this type does not offer any kind of evaluation of the likelihood of any of these scenarios emerging.

Additionally, many of the relationships between individual elements may not be explicitly defined within the actuarial models that were used to create the best estimate results. What is the impact on policyholder withdrawals when the news headlines exclaim declining birth rates and increased longevity? Perhaps this means annuity policyholders will delay withdrawals, understanding that they will need to live longer and are reminded that they have fewer children to supplement their retirement lifestyle. Or perhaps they will accelerate withdrawals because they will feel like it's better to spend their funds now rather than wait for some future that may not appear.

What is the impact on sales, and subsequent cohort mortality, if the average agent recruited to the agency force is ten years younger than in the past five years? What might be the expected market impact of venture capital companies buying several smaller competitors over the next decade? Will it drive increased demand for the product? Reduce prices? Accelerate employee turnover? Create a deluge of new applications as policyholders seek to replace current policies with other brands once they find out their insurer is now owned by a hedge fund?

Because many of these factors are not explicitly identified in the actuarial models, there may be no way to measure sensitivities of expected financial outcomes to some of these "hidden" drivers. However, a GenAI model, trained on historical company financial results, as well as external factors, could provide a more robust understanding of the various possible scenarios and their likelihood of emerging.

Suppose that an insurance company trains a GenAI model on the past 10 or 20 years' worth of company financial information, industry news and experience studies, and external factors, such as Treasury rates, Wall Street Journal articles, and Congressional Budget Office projections. Then, after creation of the best estimate results, an actuary could ask the GenAI model to also provide some potential future outcomes expected over the next few years, along with the impacts of those various predictions and the likelihood of occurrence.

In this way, the typical "sensitivity analysis" can be expanded beyond linear sensitivities to understanding not just the impacts, but the likelihood of such trends emerging.

A second step after the creation of potential scenarios is to outline the potential responses by the company to such emerging situations. That is, the GenAI model can be asked to provide specific strategies and actions in response to the emergence of various external or internal scenarios. If the GenAI model is also supplied with something like a metric for measuring the strength and weakness of those strategies, then the potential responses can also be relatively ranked against each other. This would provide the company with valuable information about the potential tradeoffs available within emergent situations.

In addition, the GenAI model may be used to quickly evaluate the impacts of certain future specific business decisions under consideration. One could ask, for example, "What if the company decides to sell two blocks of somewhat unprofitable business? How might the industry react to this news? Will it impact our parallel effort in expanding our other divisions into new markets?" Because the GenAI model can incorporate many more dimensions of interaction than typical actuarial analyses designed to be very precise, it may offer the potential to examine a significantly broader scope of potential impacts.

These exercises are less relevant to making specific predictions about which futures are expected to emerge. Rather, they are most useful for understanding the potential scope of possible futures, developing a plausible story that may lead to one of those futures emerging, and even providing some early indications that the potentials are collapsing into one specific scenario so that company leaders can begin to implement the response actions identified previously.

3.6 AUTOMATION AND EFFICIENCY

3.6.1 AUTOMATION

The most straightforward efficiency-creating application of Generative AI is to generate summaries: recent financial experience, product descriptions, market evaluations, and the like. For example, a user could train a GenAI model with the following:

- Historical financial plans at whatever level of detail is desired.
- Actual financial results covering those same time periods. These may be spreadsheets, printed or online reports, or links to financial reporting files.
- The internal discussions around how to explain the deviations of the actual experience from the planned values. This may come from chat logs, emails, or transcripts of meetings held to discuss financial results and their implications.
- The final external presentations (to company management, board members, shareholders, and regulators). These may be in the form of printed documents, online reports, or transcripts of shareholder calls.

Then, the user could ask the GenAI model to summarize trends in either the numbers or the texts explaining those values. The GenAI model may also be asked to summarize trends in deviations from historical financial plans, allowing for more in-depth synthesis of trends than could be discovered by an individual evaluating those interactions one at a time.

Additionally, an actuary could use this trained model to analyze the future forecast of financial performance, and anticipate potential deviations from this plan. As experience emerges over subsequent quarters and years, the GenAI model can be used to create the summary explanations of deviations from the plan. The actuarial user should always check these explanations to ensure that the model has not either hallucinated experience or provided nonsensical explanations for the variances, yet this can be a significant time-saving device to accelerate review and approval of financial results.

3.6.2 EFFICIENCY

Understanding and Optimizing the Chain of Data

A second application of Generative AI may be in investigating the trail of data necessary to perform specific operational tasks. Actuaries at all levels will be familiar with having several linked spreadsheets as part of a process. That is, values are first generated outside of a process and entered into spreadsheet A, where first-level outputs are calculated. These outputs are referenced via links in spreadsheet B, where second-level outputs are calculated, and so on. This chain extends potentially for many levels and potential interactions, whereby spreadsheet D may reference all of spreadsheets A, B, and C, in order to produce its desired outputs.

In order to investigate the chain of data veracity, a user would have to open each individual spreadsheet and document the upstream and downstream elements, manually tracking the increasingly-complex spider web of references. Historical experience with this process has resulted in an inventory of hundreds of upstream elements and dozens of downstream references for a single value. Understanding the implications of errors in this process and planning any changes to the number is a daunting task for an individual.

However, GenAI could be applied in this situation by asking the AI model to investigate and document the full chain of upstream and downstream links relevant to a specific spreadsheet, process, or output report. The output from the GenAI model would, therefore, be a list of all the source elements necessary to create the target output element. As well, the GenAI model can provide a list of potential downstream impacts if a specific element or report is changed, modified, or canceled outright, either accidentally or in the pursuit of operational efficiency.

Instead, that efficiency can be created by asking the GenAI model to create a streamlined approach to producing the target output, whether that be a single number or a report. The user can task the GenAI model to identify the minimum number of input items to be included and what their source is. The GenAI model can then summarize the calculation steps that may be needed in order to generate the specific output, and can also create comprehensive documentation about which downstream reports also use that specific output element.

Understanding this data transfer chain and creating a documented, as-simple-as-it-can-be process can shorten the time to create reports, minimize the potential for error, and eliminate many human tasks of copying and pasting from one file to another.

Using Generative AI to re-engineer an entire process may not seem worthwhile for a single value or report. However, expand this idea to all of the monthly or annual reports that an insurance company creates, and the potential efficiency gains soon emerge. Applying GenAI tools to the task of discovering the shortest path through the full range of repetitive tasks will significantly decrease the time to complete the whole portfolio of processes.

Efficient Reviews of Others' Work

GenAI tools could be tasked to review all of the individual formulas within spreadsheets, all of the macros hidden from the typical user, and all of the hard-coded values within reference tables, for reasons such as:

- To discover whether any hard-coded references exist and, if they do, make them visible and documented so that they are updated at the expected cadence.
- To evaluate potential hidden problems, such as coding within macros that is inconsistent with surface-level data formatting, reference ranges, and so on.

Using GenAI in this way to provide additional support for peer review will enable the actuary to focus less on the technical minutiae of whether the specific spreadsheet cells are coded correctly. Instead, actuaries will be able to perform the value-added task of evaluating whether the entire spreadsheet, program, or process produces results that are reasonable, given the specific inputs and intentions of the spreadsheet, program, or process.

3.7 CLAIMS

GenAI can augment claims processing and assessment in various ways. The most straightforward would be to use GenAI to process general insurance claims via photographic evidence. This is possible because of the new tool's ability to recognize components of images and reverse-engineer what the "prior" state would have looked like.

Suppose that an insured driver is in an accident. By supplying photos or videos of the damage, and supplementing with additional questions prompted by the GenAI assistant as to details of the accident, the GenAI-augmented process could perform the following functions:

- Confirm the make and model (and even VIN) with policy details in the administrative system, validating that the vehicle in the images is actually covered by the policy;
- Estimate the amount and type of damage to the vehicle that would be consistent with the specific accident, versus any pre-existing damage that should not be covered (reducing the chance of claim exaggeration);
- Provide an estimate of the actual repair costs;
- Connect to a list of trusted repair facilities in the area;
- Notify those facilities, providing details of the accident and the anticipated amount of work to be done, request a quote from them to perform the repair, and an estimate of any availability over the short-term time horizon;
- Provide this combination of cost and availability to the insured in order to make an informed repair decision;
- Coordinate scheduling any appointments related to the claim, such as towing or doctor's visits necessary;
- Track the progress of the claim over time until completion;
- Notify the actuaries involved if any actual-to-estimated claims deviate outside of a stated error bound, so that they can make modifications to any aggregate forecasts;
- Notify the actuaries if recent claims tend towards one end of the error bound or the other, so that they can initiate experience studies sooner than normal in order to examine whether the GenAI-perceived trends are, in fact, justified.
- Throughout, the automated assistant can communicate with the insured, informing them of progress and next steps, respond to questions, and create a documentation trail to support the decisions made along the way.

This is a simple example of how GenAI may improve the claims process for straight-forward auto insurance. The goal is to reduce the administrative burden on claims adjusters, allowing for increases in caseload management, reduction of error, and increasing customer satisfaction. Such differentials should reduce overall administrative costs, as well as allowing the actuary to more accurately project claim costs for new coverages.

Other applications of this type of process automation and augmentation with GenAI models include:

- Auditing - use the GenAI model to develop a series of questions to ask during a claim review to ensure that all relevant aspects of the company's risk management strategy are included;
- Customer insight - use the GenAI model to scan the entirety of customer communications (across email, chat, and phone records) and develop trends in complaints or questions; then, use the model to create a set of actual frequently asked questions and a portfolio of responses to such questions in different tones or styles; and
- Fraud detection - as noted above, when a claim is issued, the GenAI-augmented process can quickly determine the likelihood that all of the damage is consistent with the details of the accident specified and that the vehicle is actually covered by the policy. Additionally, because of the access, in real-time, to the large datasets of the company, the model may be able to notice anomalies in this claim versus other claims, which are supposed to be similar, and notify responsible parties that this claim should be investigated more thoroughly.

Beyond auto insurance, this type of process automation could apply for medical claims, disability, other property and casualty insurance, and even life insurance, where the insured can take a photo of a death certificate in order to start the process.

3.8 UNDERWRITING

As the purpose of underwriting is to appropriately classify the riskiness of the potential insured, it may seem like underwriting applications of Generative AI may not affect actuaries. However, this is short-sighted, as underwriting affects every area of insurance policy operations.

For Individual coverages

Individual coverages, at the lowest level, are already experiencing automation in the underwriting process, from guaranteed issue policies to accelerated underwriting processes.

These automations simplify and streamline the process by asking customers to input their information directly into administrative systems, rather than enter information on physical forms which requires an insurance company to re-key the same information. Such efficiency is the low-hanging fruit of underwriting improvements.

Once the relevant information is stored in the administration system, GenAI applications can then scan large databases, looking for consistencies within the information entered on the company-specific forms and the information available in other areas. For example, suppose that an applicant accidentally mis-keys a birth year as 1945 instead of 1954. Such an easy error would make it appear that the policyholder is nine years older than she actually is, and lead to incorrect classification. However, if the GenAI model has access to public records databases, there could be a cross-reference to any of that public information about this same individual. The GenAI application can then reply that it has found conflicting information in other places; that this individual has stated her birth year is 1954, not 1945; and then ask the individual to confirm which is the actual birth year.

Another way that GenAI may be involved is with larger coverage levels and more unique coverages. Actuaries typically base rate classifications on large volumes of similar historical data. For unique coverages, or for those coverages where the company is taking on a risk profile that may be substantially different from their typical insured, such large volumes of comparable risk profiles may not exist.

GenAI can help by synthesizing large amounts of data from internal and external sources. As these unique or different coverages may not line up exactly against historical experience, actuaries may feel as if they must add additional risk mitigations, often in the form of higher premiums and capital requirements, in order to compensate for this uncertainty. GenAI models can be trained on the hundreds or thousands of defining characteristics that these unique situations may present and provide a relative risk score of this potential policyholder versus the historical policyholder profile. The actuary could work with underwriters and the GenAI model to develop suggested lines of questioning to ask the applicant in order to appropriately understand the risks involved. This may result in an iterative, customized process that cycles through several question-and-answer feedback sessions with the prospect in order to adequately understand the various risk factors pertinent to the situation.

This would contrast with the fixed set of standard questions asked of every single applicant for which accelerated underwriting is used. A single set of questions is appropriate for smaller coverages where the risks are more easily mitigated through the law of large numbers. However, for more unique coverages where the benefit of diversification will not be realized, a customized process is necessary.

For group coverages

An additional underwriting application may be available for group coverages. Challenges herein include incomplete data and the shifting coverage demographics from period to period. Similar to the situation

above, underwriters and actuaries could work to train the GenAI model on the principles of risk mitigation and risk classification to be applied to new groups as they are added to the company's insured population. The actuary and underwriters can then request the GenAI model to:

- Evaluate the similarities and differences of the new groups to the existing population;
- Suggest a line of further questions that would help the company to understand what drives this group to have a risk profile that differs from prior populations;
- Test the validity of the answers provided, offering a ranking of the likelihood of truthfulness for each of them, based on the consistency of those answers to how other respondents answered;
- Provide an estimate of the relative risk of this group versus others, based on the similarities/differences and the truthfulness quotient; and
- Estimate the impact on the entire company's risk profile should this group be accepted for coverage.

Always up-to-date information

One additional benefit of using GenAI automation assistance is that everything is automatically connected to everything else, updating either in real-time or, perhaps, daily. This way, there is no lag between historical and current information. As well, summarized histories of these kinds of individualized risk assessments can be included as reference material for the actuary when performing experience studies, company forecasts, and capital allocation. If the actuary has a better understanding of the risk profile being covered, it is likely that the reserves allocated to individual groups or claims, and the capital allocated to specific lines of business, can be more appropriately determined.

Note, it is not necessarily that reserves or capital are intended to be reduced, either individually or in aggregate. The point is to allocate reserves and capital in a manner that accurately reflects the best estimates of the risks involved, resulting in better decision-making regarding accepting and rating individual cases or an entire block. This should enable more accurate projections of future income statement and balance sheet values as risks evolve and the portfolio matures.

Section 4: Practical Application of Generative AI

In the previous sections, we've described some of the strengths of Generative AI and potential applications for actuaries. In this section, we consider how to go about harnessing the potential of Generative AI.

However - we start with some further exploration of the potential limitations and areas of concern.

4.1 LIMITATIONS AND AREAS OF CONCERN WHEN USING AI

Quality and Accuracy - Generative AI can be prone to 'hallucinations' - plausible sounding, but inaccurate, outputs. Some AI platforms are investing heavily in their tools and algorithms to mitigate this so the choice of model is important. However, it remains a concern and users should ensure that any outputs from Generative AI models are tested and/or checked appropriately.

Repeatability and Drift: Drift refers to the phenomenon that, over time, Generative AI models can provide very different outputs even given the same inputs. This can be a function of both the elements of randomness in the underlying algorithms and the evolution of the underlying training data. It can present a challenge for the use of AI in ongoing processes where there is an appetite for stable, reproducible outcomes. This can be controlled by implementing guardrails in the way that prompts and additional data are implemented.

Data Privacy and Security - Ensuring that AI models comply with privacy laws and maintain confidentiality. This presents something of a conundrum for companies - the need to balance an appetite to access the potential gains and insights that Generative AI can offer versus the need to keep data confidential. This is doubly so for insurers given the private and sensitive nature of the data they hold. This can be addressed through the choice of AI model/platforms and the deployment model, for example, by being deployed onto private networks where data is not transferred.

Auditability and Transparency – The complexity of Generative AI (or machine learning) algorithms presents a challenge for users when it comes to understanding how an algorithm came to a particular outcome. This is especially the case where these algorithms make decisions that affect policyholders, such as in processing claims or underwriting decisions, which may be subject to challenge or where the need to demonstrate fairness or lack of discrimination is paramount.

Copyright - Generative AI models are often trained using significant amounts of data (text, code or images) publicly available via the internet. Other models are trained on private sources of data. Those models trained on public information raise questions around the impact of training on copyrighted data. There is also the risk of potential infringement by AI - for example, if it produces outputs that are very similar to copyrighted works. There is the question of the ownership of outputs of Generative AI - is it the prompt engineer or the owner of the platform who controls the rights to any intellectual property created using such tools?

Quality of Input and Training Data - Actuaries are familiar with the premise that any model is only as good as the inputs. In the public datasets upon which many Generative AI models will be trained, only a very small proportion of the content in that training data will be actuarial in nature. This should give actuaries pause when considering generic GenAI for actuarial applications, as the models may not be as appropriate as for generating other types of outputs.

Fraudulent Use - The ability of Generative AI to create realistic images, text and documents has the potential to support fraudulent or other nefarious activities. For example, creating altered images that show greater-than-actually-happened damage to cars or properties, and subsequently submitted as part of

a claims process, could lead insurance companies to pay additional claims than truly occurred. Altered or faked documents could also lead to incorrect risk assessments during underwriting, leading to incorrect classifications, insufficient premiums, and inappropriate capital allocations.

Prompt Engineering - Prompt engineering is the art and science of carefully crafting the input prompt to return a useful output. A key issue with prompt engineering is that it may require a substantial learning curve. Unskilled or careless prompting may end up producing results that are less valuable than what would emerge using skilled or careful prompting. Additionally, as some GenAI models “learn” from the prompts fed to them and the subsequent fine-tuning of follow-up prompts, careless prompt engineering could introduce additional biases into the models.

Ethical and Moral Considerations

- Bias - AI models are trained on available data, but to the extent that real world data may contain biases, there is a risk the AI systems will reflect or exacerbate any bias in the data used to train models.
- Transparency - of the model’s “inner workings” is a major concern that has not been fully satisfied, as many of these AI models have billions or trillions of parameters underlying them, and some of the relationships within being undefinable using typical algorithmic methodologies. This makes it difficult to understand why some AI models produce the results they do and can lead to irreproducibility (where similar inputs produce different outputs).
- Impact on jobs - the potential for using GenAI tools to replace some job functions, like crafting product descriptions or creation of marketing materials, leads to ethical questions about the impact on those who are replaced or pushed out of roles as the use of GenAI advances.
- Accountability - the issue about accountability for results will always remain. Users may deflect responsibility for content created over to the tools rather than themselves in an effort to avoid negative implications of the use of those outputs. Allowing poor use of any tool should not be accepted. However, in the future, it may be more prevalent to attempt to excuse bad outcomes as simply a function of the model or the AI, rather than the responsibility of those who decided to use them.
- There is increasing interest in the political sphere in the regulation of AI platforms and services. As with companies, lawmakers are trying to balance the potential gains in economic productivity and growth with concerns around the impact on jobs and a general sense of the technology evolving faster than society’s ability to process the change. For example, the European Union published its first [AI regulation act](#) in 2023.

The good news is that most of these concerns can be addressed - or at least mitigated - with appropriate choices of deployment architectures, choice of foundational models, and training and fine tuning the approach.

4.2 ARCHITECTURAL AND DEPLOYMENT CONSIDERATIONS

While manually typing prompts into a Generative AI tool may be manageable for basic queries and some ad hoc tasks, it is unlikely to deliver the benefits at a scale that will drive meaningful change for insurers. Similarly, uploading sensitive data to public AI platforms is likely to, at a minimum, fall foul of internal data policies and could be a major breach of data privacy laws.

As such, it is important to think about how Generative AI can be plugged into the broader ecosystem of data, applications and use cases across an insurer. Success will require a strategic approach with buy-in from corporate IT, Data, Information Security teams, as well as individual business lines.

These considerations will drive decisions around which models, platforms or applications to use, as well as the training data and prompting mechanisms required to deliver the desired outcomes.

4.2.1 MODELS AND PLATFORMS

There is an extensive and growing list of Generative AI tools available in the market. These often vary by specialization, such as writing code, creating narrative or expository content or producing and interpreting images. They may also vary as to whether they are a ‘platform’ that link to one or more underlying Generative AI models (sometimes called foundation models) that can be adapted to multiple use cases, or whether they are AI-powered applications with a particular specialism.

For example, OpenAI offers the ChatGPT platform. The platform is equipped with an API that allows users to structure solutions to their particular needs. It also offers access to different models, such as GPT-3.5 or GPT-4, that have different characteristics in terms of the number of parameters and the scope of the training data. ChatGPT users can access a marketplace of AI-powered applications that specialize in different tasks. It is worth noting that sometimes these distinctions are clear but, in other cases, the dividing line between the model and platform may be less obvious.

Another significant point of differentiation between models and platforms is in their approach to training data and the treatment of new input data. Some will be trained on widely available public data, while others can be trained exclusively on internal company data - with many variations and hybrid options.

The choice of whether to leverage a platform, build an ecosystem around a foundational model, or use the functionality of a specialist AI-powered application will very much depend on the task at hand. Other factors will be how users will interact with it, and the task-specific and security-related data issues. We illustrate some examples below.

Table 1

GENERATIVE AI PLATFORMS: A NON-EXHAUSTIVE LIST OF GENAI PLATFORMS AT THE TIME OF WRITING

Vendor	Platform	About
OpenAI	ChatGPT	LLM platform, text-based responses to input prompts. Supports an evolving series of models such as GTP-3.5 and GPT-4. Also offers a suite of third-party plugins with specialist capabilities.
Amazon Web Services (AWS)	Bedrock	Cloud-based managed Service platforms that support a range of proprietary and third-party GenAI models and integration options. Integration with broader ML tools.
Microsoft	Azure AI studio	
Google	Google AI studio	

Generative AI (Foundation) Models: Overview of Generative AI Models

There is something of an arms race in the capability of the underlying AI models with new players entering the market and existing vendors offering a stream of upgraded capabilities. While there are various benchmarks for performance in a given task, such as lines of code created or volume of text generated, much of the performance of a model is based on the number of ‘parameters’ - indicative of the volume of data used to train the model and ‘context length’ - the size of the allowable input prompt measure in ‘tokens’ (words or parts of words).

Whether bigger is better very much depends on the context. General purpose chatbots typically have very large parameter counts reflecting the size and diversity of the datasets. While not officially disclosed, some estimates put the number of parameters in OpenAI’s GPT model at over 100 trillion. The downside of these

large models is the requirement for vast amounts of computing power and they are typically only accessible through cloud computing platforms.

Some models, such as [Meta's Llama 2](#), are available in different sizes. This allows more flexible deployment options - for example, within a companies' own IT infrastructure or even onto individual devices.

Another important consideration is the extent to which the models include mechanisms to limit quality issues. Many models include a feedback loop that allow the outputs to be refined by both automated and human feedback (Reinforcement Learning from Human Feedback or RLHF) – see, for example, this article from Anthropic about their [Claude models](#).

The rate of change in these models, as well as their individual specializations, makes ranking individual models challenging, though a search for 'best AI models' certainly provides plenty of results. For example:

- [TechTarget](#) provide a description of some common models with details about specialization and number of parameters
- [Simplilearn](#) describe a range of popular models and tools and opine on pros and cons
- [Stanford University](#) have constructed an index that attempts to measure the transparency of various AI models

Table 2

AI-POWERED APPLICATIONS: A SMALL SELECTION OF AI-POWERED APPLICATIONS REFERENCED IN THIS PAPER

GitHub Copilot	Coding/software development productivity tool. Assists developers by suggesting code and entire functions in real-time within the coding environment.
Amazon Code Whisperer	Coding/software development productivity tool. Assists developers by suggesting code and entire functions in real-time within the coding environment.
Microsoft Copilot	General productivity. Helps automate and assist with workflow tasks within the Microsoft 365 suite.
Julius.ai	AI-powered data manipulation and analysis tool.
Dall-E	Generates original images and art from textual descriptions using AI image generation technology.
Copy.AI	Marketing copy and content, such as blog posts and social media updates, using AI.
Fireflies.ai	Automates note-taking and meeting summaries from voice conversations in real-time.

The above is only a small subset of the available options with the list constantly growing and the options continually improving. There are several trends in the evolution of these AI models and platforms:

- Improving robustness and accuracy by incorporating feedback loops and addressing some of the potential bias issues that may arise.
- Improving the efficiency of models in terms of the computing power required, which has both environmental and cost implications.
- Improving versatility - supporting a wider range of applications.
- Improving the range of deployment models.

When considering the use of any AI model, platform or tool, it is important to research the available options at the time.

4.2.2 TRAINING PROCESS AND DATA

Actuaries are familiar with the principle that models are only as good as the data that underlies them, and this is equally true for Generative AI models. The quality, quantity, and diversity of data directly influence the effectiveness and capabilities of these models.

Generative models require a significant amount of data and computational power for training. The models learn to generate new content by adjusting their internal parameters to reduce the difference between the generated output and the real data. Exact details on training datasets for proprietary models are often scarce but, as an example, OpenAI's GPT-3 model data is described as including hundreds of billions of words and more than 45 terabytes of text input data.

For Generative AI models to produce accurate and realistic outputs, the input data must be of high quality. This means data should be clean, well-organized, and free from errors and inconsistencies. It's also crucial that the training data is representative of the real-world scenarios where the model will be applied. Lack of representativeness can lead to biased or skewed results.

The volume of data contributes directly to the model's ability to learn and generalize from complex patterns. That is, more data will lead to more robust model capabilities and greater application to various use cases. The training data should encompass a wide range of examples to cover the different variations the model might encounter. This is crucial for the model's ability to generate diverse and novel outputs. Diverse data helps in avoiding overfitting, where a model performs well on training data but poorly on unseen data. These volume requirements must be considered in relation to the capabilities of the underlying foundational model. A nano model with a few billion parameters requires much less training data than a high-end model with hundreds of trillions of parameters.

Different applications also have specific data requirements. For instance, image processing applications (such as a medical image or damage assessments in claims processing) require high-resolution, task-relevant images, while natural language processing relies on a wide variety of textual data.

The way data is sourced and used is subject to ethical considerations. Issues like user consent, privacy, and data ownership are crucial. For actuaries, relevant data may contain sensitive policyholder or health information, which is subject to stringent legal and regulatory requirements.

Given these considerations, a key question is whether a pretrained model can be used to perform the task, or whether the model will be trained as part of the overall solution design. There are several mechanisms by which pretrained models can be leveraged, but tailored to more specialist tasks, and we explore these in the next section.

4.2.3 PROMPT ENGINEERING, TRANSFER LEARNING AND RAG

A key part of using Generative AI is how you specify what you want it to do and what additional information is provided to tailor the outputs to the task at hand.

Prompt Engineering

Prompt engineering refers to the practice of carefully crafting input prompts to guide an LLM's responses in a particular direction or to fulfill specific criteria. It's important because these models respond to input prompts, and the structure, wording, and content of those prompts can significantly impact the output.

Important considerations for prompts include:

- **Specificity:** By making the prompt more specific, a user can guide the model to produce more detailed and relevant answers.
- **Context:** Including the right context helps the model understand the domain or subject matter more clearly.
- **Constraints:** Adding constraints to control the style, tone, or format of the output makes that output more appropriate to the requested use.

One of the key advantages of prompt engineering is that it allows the user to "specialize" the model for various tasks, without additional training, by using carefully chosen prompts. This is especially true for LLMs with very large training datasets where users may adapt the model to different requirements by merely changing the prompts, giving flexibility across applications, i.e., you can use the same LLM to summarize documents, generate python code to reformat a data file and give you a lesson in the operation of Generative AI.

That flexibility does come with some challenges:

- **Complexity:** Crafting effective prompts might require experimentation and a deep understanding of both the task at hand and the model's capabilities.
- **Sensitivity:** Small changes in the input prompt can result in significant differences in the outputs. This could be a particular limitation if trying to use Generative AI to perform repetitive tasks.
- **Bias and Ethical Considerations:** Uncritical (or malicious) prompt design can introduce biases or lead to outputs that may not align with ethical or societal norms.

Prompt engineering is a valuable skill when working with large language models, offering a way to guide and control the model's output effectively. It's akin to asking the right question in the right way, allowing for more precise, relevant, and contextually aware responses. As such, prompt engineering requires a combination of understanding both the model and the task it is being used to perform, and often some trial and refinement to produce output of sufficient quality.

The following is an example of how different prompts, related to explaining a Universal life policy, result in different output: <https://chat.openai.com/share/189a4571-3c0f-4661-9753-af2ed53d01c8>

Transfer Learning

Transfer Learning is primarily used to leverage the knowledge gained from one task to improve performance on a different, but related, task.

It often starts with a model that has been pre-trained on a large, general dataset. The pre-trained model is then fine-tuned on a smaller, specific dataset related to the new task. The necessity for transfer learning depends on the application. For instance, if a new task requires a deep understanding of a specific domain not covered in the original training data, transfer learning is one way that this additional data may be included in the models capabilities.

A key benefit of transfer learning is that it allows for efficient model training with less data and fewer computational resources, as the model doesn't need to learn everything from scratch. It is common in many areas of AI and machine learning.

Both prompt engineering and transfer learning rely on the use of pre-trained models. While transfer learning adapts models to new tasks through additional training, prompt engineering leverages their existing capabilities through careful crafting of the prompts without the overhead of additional training.

Transfer learning and prompt engineering are typically complementary. A model fine-tuned on a specific task through transfer learning might still benefit from carefully engineered prompts to perform optimally, especially in the case of language models.

Retrieval Augmented Generation

Retrieval Augmented Generation (RAG) is another advanced technique in the field of Generative AI that supplements the capabilities of a pre-trained LLM with additional, specific data to enhance the outputs of the model with task-specific content.

How RAG works:

- **Retrieval System:** At its core, RAG uses a retrieval system (like a search engine) to find relevant documents or information based on the input query. This system can access a vast database of text or can be tailored to point to specific reference documents.
- **Integration with Language Model:** The retrieved documents are then integrated with the outputs of a language model. This model, often a variant of a transformer-based model, generates responses based on both the input query and the information from the retrieved documents.
- **Enhanced Context and Information:** By combining retrieved external information with its pre-trained knowledge, the language model can provide more accurate, detailed, and contextually relevant responses.
- **Refinement and Output:** The language model refines the information and generates a response to the query.

RAG-based systems are likely to be a common deployment model for commercial applications as it allows the system to be more easily tuned to company-specific data and use cases.

Comparisons between Transfer Learning and RAG

Each of these strategies relies on the use of pre-trained models, but they incorporate task-specific data in different ways:

- **Transfer Learning** is about transferring knowledge from one domain to another within a model, while **RAG** is about augmenting a model's generation process with external information at the point of query.
- **Model Training vs. Model Execution:** Transfer Learning is a strategy used during the training phase of a model. In contrast, RAG is employed during the execution or inference phase where the model interacts with external data sources in real-time.
- **Data Dependency:** Transfer Learning reduces the dependency on large amounts of task-specific data to fully train a model. RAG relies on external data sources to provide additional context during generation.
- **Internal vs. External Knowledge:** Transfer Learning enhances a model's performance based on internal knowledge (what the model has learned). RAG extends a model's capabilities with external knowledge (information outside the model's training data).

Chaining

For complex tasks, it may not be possible to complete the task in a single step. Chaining involves sequentially using multiple AI models or processes where the output of one serves as the input for the

next. This is common in complex tasks where no single model can handle the entire process end-to-end, and it's likely to be a key requirement for practical applications of Generative AI for actuaries.

Evolving Methodologies

The training and prompting of an AI model to perform a task is a critical step. It is important to consider the information and context the model requires to perform the task. Both Transfer Learning and RAG aim to enhance the performance and capabilities of pretrained models beyond simple prompt engineering, but they do so in different stages (training vs. execution) and through different mechanisms (knowledge transfer within a model vs. augmenting generation with external data).

This is also a rapidly evolving field with new techniques and practices. Publications such as [Medium](#), [TowardsDataScience](#) and [TowardsAI](#) regularly include articles on these evolving methodologies.

4.3 CHECKLIST FOR USING GENERATIVE AI

As discussed in the previous sections, there is a wide range of options when it comes to model choice, deployment architecture, training data, and prompt engineering guidelines. These choices will directly affect the quality of outcomes from using Generative AI. Some questions to consider as an organization prepares to enhance their business processes with Generative AI:

About the task:

- What is the task?
- What is the specific answer which is to be provided? How will that answer be used?
- What data and context are required to perform the task?
- How many end users will access the solution?
- How will they interact with it?

About the solution:

- Is there a third-party platform or application that performs the task or does a customized solution need to be developed?
- Can a pretrained model perform the task? Or does a similar model have to be trained for the specific task?
- Is prompt engineering sufficient for the model to perform the task?
- Does it need to be tailored to the specific task with transfer learning or RAG?

About the data:

- If the model is pretrained, does it include sufficient training data relevant for the usage?
- Is sensitive data (perhaps subject to greater security, such as personally-identifiable information) needed for additional context prompting?
- Will sensitive data leave the corporate network?
- If outside the network, does it become part of the model's training data (which could risk exposure)?

About quality requirements:

- How sensitive is the organization to the quality of outputs?
- What are the risks of erroneous outcomes?
- What processes are in place to ensure review and quality of output?

Commercial considerations:

- Is the model or technology appropriately licensed for commercial use?
- Have costs been accounted for? Typical costs include token-level costs for API calls, or per-user or enterprise licenses for vendor hosted solutions; or up-front infrastructure, maintenance, and energy costs for internally deployed solutions.

Organization readiness:

- Have relevant stakeholders provided the necessary approvals - management, Information Technology, Information Security?
- Are there any ethical concerns about using a commercially-available solution in this way?
- Do users and stakeholders understand the risks and limitations?
- What governance processes are required?
- Can procurement processes support the acquisition of the relevant technologies?

Section 5: Summary (Conclusions)

Artificial intelligence (AI) is the current frontier of technology across not only the actuarial profession, but the insurance industry and the general economy. Generative AI, where the application is intended to generate some additional output based on a large volume of input data and a defined prompt, is likely to support actuarial work in substantial ways.


Actuaries may find themselves using Generative AI in a variety of applications:

- Understand model code;
- Develop large volumes of synthetic data to test models or infrastructure capacities;
- Create automated workflows that seamlessly process claims;
- Perform model documentation and create an audit trail that supports a robust model governance process;
- Learn quickly how to interpret or develop code for new languages;
- Design targeted training programs for new employees or junior staff so that learnings can be customized to the individual; and
- Quickly summarize results of regular analyses and identify trends or anomalies.

In addition, actuaries may find their work impacted by the use of Generative AI elsewhere in the organization. This could include marketing and product development shifting the average profiles of agents, employees, and customers; or asset/liability management decisions influenced by the results of generated scenarios trained on proprietary information about the company.


The front-end applications that expose the underlying models themselves vary, from stand-alone browser-based apps to co-pilots “looking over the shoulder” of professionals and supporting their progress. These end-user applications will support a range of use cases, while the underlying Generative AI models will have their own range of applications. When deciding whether to introduce a Generative AI application, care should be taken to consider the technical, operational, practical, commercial, ethical, and regulatory aspects of the application, the underlying model, and any outputs created from them.

The commercial landscape is changing rapidly, affecting actuaries in all segments. This paper presents a point-in-time overview of many substantial themes, yet is not intended as any kind of standard of practice or best practice. By understanding the opportunities and implications of including Generative AI in their business, actuaries are poised to support substantial change and growth of their companies, the broader industry, and the profession.



Give us your feedback!
Take a short survey on this report.

[Click Here](#)



Section 6: Acknowledgments

The researchers' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the Project Oversight Group for their diligent work overseeing, reviewing and editing this report for accuracy and relevance.

Project Oversight Group members:

Henry Chen, FSA, MAAA, FCIA

Bruce Friedland, FSA, MAAA

Michael Levine, FSA

Mark Ma, FSA, MAAA

Paula Schwinn, FSA, MAAA

Feng Sun, FSA, CERA

At the Society of Actuaries Research Institute:

Korrel Crawford, Senior Research Administrator

R. Dale Hall, FSA, MAAA, CERA, Managing Director of Research

Appendix A: Summary of Key GenAI Models and Techniques

The inner workings of Generative AI is a complex area with a great deal of technical complexity. We introduce some of the key concepts below, though we omit much of the detail.

Some additional primers on Generative AI:

- <https://www.simplilearn.com/tutorials/artificial-intelligence-tutorial/what-is-generative-ai>
- <https://medium.com/@ramki.desiraju/cxos-primer-to-generative-ai-foundation-models-from-autocomplete-to-autocomplete-a3817943656e>

NEURAL NETWORKS AND DEEP LEARNING

Generative AI predominantly relies on neural networks, specifically deep learning models, to understand and generate new content. These networks are composed of layers of nodes or neurons, interconnected with weights that are adjusted during training. Deep learning models can have hundreds of layers and millions of neurons, enabling them to capture complex patterns and relationships in data.

- **Convolutional Neural Networks (CNNs):** CNNs are primarily used in image generation and processing tasks. Their structure is well-suited for detecting patterns and features in images, making them ideal for tasks like image generation, enhancement, and style transfer.
- **Recurrent Neural Networks (RNNs) and Variants:** RNNs are designed to work with sequential data, such as text or time-series data. They have the unique feature of retaining information from previous inputs through internal memory. RNNs, especially LSTM networks, are widely used in generating coherent and contextually relevant text in models like language models and chatbots.

TRANSFORMER MODELS

Transformer models have revolutionized text generation. Unlike RNNs, they process input data in parallel and rely heavily on attention mechanisms. Attention mechanisms are the transformer's superpower. It looks at each word (or piece of data) in the context of all the other words in the sentence. For example, in the sentence "The cat sat on the mat," when it looks at the word "sat," it also considers "The cat" and "on the mat" to understand its full meaning. This is done through a process called 'self-attention,' which allows the model to weigh the importance of all other words when understanding each specific word.

Thus, a key advantage of transformers is their ability to look at the whole sequence of data (like a whole sentence) at once, rather than one piece at a time. This makes them good at understanding context and meaning in language processing tasks.

Another major benefit is their parallel processing capability. Transformer models can be scaled up efficiently, leading to allowing them to process large volumes of inputs and outputs in tasks like language translation or text generation.

VARIATIONAL AUTOENCODERS (VAES)

Variational Autoencoders (VAEs) (see Kingma & Welling, 2019) are a type of generative model that are particularly effective in producing complex data like images, music, and text. To understand VAEs, imagine them as expert data compressors and reconstructors.

VAEs consist of two parts – an encoder and a decoder. The encoder compresses the data into a simpler, condensed form (latent space), and the decoder reconstructs the original data from this compressed form.

Latent Space Representation: The latent space is a lower-dimensional space where the data's most crucial characteristics are stored. This space is what the VAE learns to manipulate in order to generate new data.

VAEs work using the following steps:

- **Encoding:** The encoder takes input data (like an image or text) and converts it into a set of parameters in the latent space. These parameters typically represent two things: the mean and variance that define a probability distribution.
- **Sampling:** VAEs randomly sample points from this distribution to generate new data points. This randomness is key to the VAEs' ability to generate diverse new outputs.
- **Decoding:** The decoder then takes these sampled points and reconstructs them back into data. If the VAE is trained on images, the decoder would generate new images that resemble the training data.

Some common applications include:

- **Image Generation and Reconstruction:** VAEs can generate new images or reconstruct missing parts of images.
- **Anomaly Detection:** By learning the normal distribution of data, VAEs can identify data points that do not fit that distribution.
- **Synthetic Data Generation:** Similarly, learning the distribution of the data allows the VAE to generate new data samples with the same characteristics.
- **Dimensionality Reduction:** Similar to Principal Components Analysis, but with a probabilistic twist.

GENERATIVE ADVERSARIAL NETWORKS (GANs)

Generative Adversarial Networks (GANs) (Goodfellow, et al, 2014) represent a significant development in the field of Generative AI. GANs have revolutionized the way machines generate realistic images, videos, and other forms of media.

A GAN consists of two neural networks, the generator and the discriminator, which are trained simultaneously in a competitive manner.

- **Generator:** The generator network takes random noise as input and generates data (like images).
- **Discriminator:** The discriminator network tries to distinguish between real data (from the training set) and fake data (created by the generator).

Adversarial Training: The generator and discriminator are trained in opposition to each other. The generator aims to produce data so realistic that the discriminator can't tell it apart from actual data, while the discriminator continuously improves at distinguishing real from fake.

Some applications of GANs:

- **Image and Video Generation:** From creating realistic human faces to synthesizing novel images of objects, GANs have wide applications in visual media.
- **Data Augmentation:** In cases where data is scarce, GANs can generate additional data for training other models.
- **Style Transfer:** GANs can modify images to mimic a particular artistic style or combine features of multiple images.

- Super-Resolution: Enhancing the resolution of images, a process known as super-resolution, is another application area.

Appendix B: Glossary of Terms

This glossary was largely created by GPT-4 using the prompt:

Can you create a glossary of terms to be included in a report on Generative AI? The scope should cover data science, machine learning and Generative AI, including underlying models and techniques. The glossary should also include any key relevant computing and technology terms.

Refined by:

Can you expand on the above including descriptions for each section, a section on the broader types of AI? Can you extend fourth section to include VAEs, as well as expanding it to include all the terms relevant to VAEs, GANs and transformers - for example, latent space, encoder and decoder, generator and discriminator, attention mechanism etc.?

Then:

Can you add a section on interacting with Generative AI tools; this should include: prompt engineering, transfer learning, RAG, chaining and APIs - but should include any and all other relevant terms?

GENERATIVE AI

Generative AI refers to the subset of artificial intelligence focused on creating new content, data, or solutions by learning from existing datasets. This includes generating realistic images, videos, text, and voice outputs. It's particularly notable for its applications in creative and design fields.

- **Large Language Models (LLMs):** Large Language Models are a type of Generative AI that specializes in understanding and generating human language. These models, such as GPT (Generative Pre-trained Transformer), are trained on vast amounts of text data and can perform a variety of language tasks, including translation, summarization, question answering, and creative writing. LLMs are known for their ability to generate coherent and contextually relevant text based on given prompts, making them powerful tools for a wide range of applications in natural language processing.
- **Natural Language Processing (NLP):** NLP is a field at the intersection of computer science, artificial intelligence, and linguistics. It involves the use of algorithms and models to enable computers to understand, interpret, and generate human language. NLP encompasses a range of techniques and tools for syntactic and semantic analysis of text, and it's fundamental in applications like language translation, sentiment analysis, chatbots, and voice recognition systems. Generative AI plays a crucial role in NLP by providing models that can generate and manipulate language in human-like ways.
- **Deepfakes:** Synthetic media in which a person in an existing image or video is replaced with someone else's likeness using neural networks.
- **Text-to-Image Generation:** AI models that create images from textual descriptions, typically using techniques like GANs or transformers.
- **Style Transfer:** The technique of applying the style of one image to the content of another, often used in image generation.

DATA SCIENCE

- **Data Science:** combines statistical methods, data analysis, and machine learning to understand and analyze actual phenomena with data. It involves techniques and theories drawn from many fields within the context of mathematics, statistics, computer science, domain knowledge, and information science.
- **Data Mining:** The process of discovering patterns and knowledge from large amounts of data.
- **Big Data:** Extremely large datasets that may be analyzed computationally to reveal patterns, trends, and associations.
- **Data Visualization:** The graphical representation of information and data to provide an accessible way to see and understand trends, outliers, and patterns.
- **Predictive Analytics:** The use of data, statistical algorithms, and machine learning techniques to identify the likelihood of future outcomes based on historical data.

MACHINE LEARNING

Machine Learning is a branch of AI and computer science which focuses on the use of data and algorithms to imitate the way that humans learn, gradually improving its accuracy. It includes various types of learning methods and algorithms for different purposes.

- **Supervised Learning:** A type of machine learning where the model is trained on labeled data.
- **Unsupervised Learning:** Learning from test data that has not been labeled, classified, or categorized.
- **Reinforcement Learning:** A type of machine learning where an agent learns to behave in an environment by performing actions and seeing the results.
- **Neural Networks:** Computing systems vaguely inspired by the biological neural networks that constitute animal brains.
- **Deep Learning:** A subset of machine learning in artificial intelligence that has networks capable of learning unsupervised from data that is unstructured or unlabeled.

UNDERLYING MODELS AND TECHNIQUES

This section covers specific models and techniques that are foundational to the functioning of advanced AI systems, particularly in the realm of Generative AI and machine learning.

- **Transformers:** A type of deep learning model that has shown great success in understanding and generating human language.
 - **Attention Mechanism:** A process in neural networks that allows the model to focus on different parts of its input for generating a prediction.
- **Convolutional Neural Networks (CNNs):** Deep learning algorithms that can take in an input image, assign importance to various aspects/objects in the image, and differentiate one from the other.
- **Recurrent Neural Networks (RNNs):** A class of neural networks where connections between nodes form a directed graph along a temporal sequence, allowing it to exhibit temporal dynamic behavior.
- **Backpropagation:** The fundamental algorithm used for training neural networks.
- **Autoencoders:** A type of artificial neural network used to learn efficient codings of unlabeled data.
- **Variational Autoencoders (VAEs)**
 - **Latent Space:** The compressed representation of the input data learned by parts of the model like autoencoders and VAEs.

- Encoder: Part of the model that compresses the input into a smaller, dense representation (latent space).
- Decoder: The component that reconstructs the input data from the latent space representation.
- Generative Adversarial Networks (GANs): A class of machine learning frameworks where two neural networks compete with each other in a game.
 - Generator: In GANs, the network that generates new data instances.
 - Discriminator: In GANs, the network that evaluates the authenticity of data instances.

COMPUTING AND TECHNOLOGY TERMS

- Cloud Computing: The delivery of different services through the Internet, including data storage, servers, databases, networking, and software.
- Quantum Computing: A type of computing that takes advantage of quantum phenomena like superposition and quantum entanglement.
- Edge Computing: A distributed computing paradigm that brings computation and data storage closer to the location where it is needed.
- Blockchain: A system of recording information in a way that makes it difficult or impossible to change, hack, or cheat the system.
- Internet of Things (IoT): The interconnection via the Internet of computing devices embedded in everyday objects, enabling them to send and receive data.

BROADER TYPES OF AI

Artificial Intelligence can be broadly categorized into several types based on functionality, capabilities, and the level of human intervention required.

- Narrow AI: AI systems that are designed and trained for a specific task. Virtual personal assistants, such as Apple's Siri, are a form of narrow AI.
- General AI: Hypothetical AI that has the ability to understand, learn, and apply its intelligence broadly and flexibly, similar to a human's cognitive abilities.
- Strong AI: Another term for General AI, emphasizing AI with human-level cognitive abilities.
- Weak AI: AI focused on one narrow task, without consciousness, sentience, or mind, also known as Narrow AI.
- Reactive Machines: Basic AI systems that operate based on present data, reacting to current scenarios. IBM's Deep Blue, which played chess, is an example.
- Limited Memory AI: AI systems that can use past experiences to inform future decisions. Most current AI systems, like self-driving cars, fall into this category.
- Theory of Mind AI: A future class of AI that would be capable of understanding and remembering emotions, beliefs, and needs that affect human behavior.

INTERACTING WITH GENERATIVE AI TOOLS

This section delves into the methodologies and techniques used in the interaction with and utilization of Generative AI tools. These methods are crucial for effectively leveraging AI capabilities in various applications, from natural language processing to image generation.

- **Prompt Engineering:** The process of crafting inputs or prompts in a way that guides the AI to produce the desired output. This is especially crucial in systems like GPT (Generative Pre-trained Transformer), where the output is highly dependent on the input prompt's structure and content.
- **Transfer Learning:** A machine learning method where a model developed for a task is reused as the starting point for a model on a second task. It's particularly useful in situations where the dataset for the second task is too small to train a full-scale model from scratch.
- **Retrieval-Augmented Generation (RAG):** A technique that combines pre-trained language models with a retrieval system. The model retrieves documents and uses them as context for generating responses, enhancing the quality and relevance of the output.
- **Chaining:** Involves sequentially using multiple AI models or processes where the output of one serves as the input for the next. This is common in complex tasks where no single model can handle the entire process end-to-end.
- **APIs (Application Programming Interfaces):** APIs allow different software applications to communicate with each other. In the context of Generative AI, APIs enable developers to access advanced AI models and functionalities provided by various platforms without needing to develop the models from scratch.
- **Fine-Tuning:** The process of taking a pre-trained model and further training it on a more specific dataset. This allows the model to become more specialized in a particular domain or task.
- **Model Inversion:** A technique used to reconstruct input data from model outputs. It's particularly relevant in understanding and mitigating privacy concerns in Generative AI.
- **Hyperparameter Optimization:** The process of tuning the parameters of a machine learning model to optimize its performance. This is crucial in ensuring that generative models produce high-quality and accurate outputs.
- **Interactive Learning:** A process where the model learns from real-time feedback provided by users. This can be seen in AI systems that adapt based on user interactions and preferences.
- **Data Augmentation:** The technique of increasing the size and diversity of training datasets by creating modified versions of the data. This is essential in improving the robustness and accuracy of generative models.

References

Actuarial Association of Europe; *What Should an Actuary Know About Artificial Intelligence? AAE discussion paper*; 2024; <https://actuary.eu/wp-content/uploads/2024/01/What-should-an-actuary-know-about-Artificial-Intelligence.pdf>

Anthropic AI; *Model Card and Evaluations for Claude Models*, 2023; https://cdn.sanity.io/files/4zrzovbb/website/5c49cc247484cecf107c699baf29250302e5da70.pdf?dl=ModelCardClaudev2_with_appendix_v1.pdf

Barenkamp, Rebstadt, and Thomas; *Applications of AI in Classical Software Engineering*, 2020; <https://aiperspectives.springeropen.com/articles/10.1186/s42467-020-00005-4>

Cuomo, Jerry; *Quantifying the Productivity Gains of Generative AI for Developers*, 2023; <https://medium.com/@JerryCuomo/quantifying-the-productivity-gains-of-generative-ai-for-developers-86fa3a76b63c>

Desiraju, Ramki; *CXO's Primer to Generative AI & Foundation Models: From Autocomplete to Autocomplete*, 2023; <https://medium.com/@ramki.desiraju/cxos-primer-to-generative-ai-foundation-models-from-autocomplete-to-autocomplete-a3817943656e>

European Parliament News release; *EU AI Act: first regulation on artificial intelligence*, 2023; <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>

Finkenstadt, Eapen, Sotiriadis, and Guinto; *Use GenAI to Improve Scenario Planning*, 2023; <https://hbr.org/2023/11/use-genai-to-improve-scenario-planning>

Government Actuary's Department, The Financial Reporting Council Limited; *The Use of Artificial Intelligence and Machine Learning in UK Actuarial Work*, 2023; https://media.frc.org.uk/documents/Research_on_the_use_of_Artificial_Intelligence_and_Machine_Learning_in_UK_actuarial_work_AK5H1We.pdf

Gurtcheff, Jeff; *5 Ways Generative AI Will Transform Claims*, 2023; <https://www.insurancethoughtleadership.com/ai-machine-learning/5-ways-generative-ai-will-transform-claims>

Harel, Katza, Marron, and Szekely; *On Augmenting Scenario-Based Modeling with Generative AI*, 2024; <https://arxiv.org/pdf/2401.02245.pdf>

Hyperexponential; *Underwriting 3.0: Explore the future of Underwriting*, 2023; <https://www.hyperexponential.com/blog/underwriting-3-0-exploring-the-future-of-underwriting/>

Lutkevich, Ben; *16 of the Best Large Language Models*, 2023; <https://www.techtarget.com/whatis/feature/12-of-the-best-large-language-models>

McKinsey & Company; *The Promise and Challenge of the Age of Artificial Intelligence*, 2018; <https://www.mckinsey.com/~media/McKinsey/Featured%20Insights/Artificial%20Intelligence/The%20promise%20and%20challenge%20of%20the%20age%20of%20artificial%20intelligence/MGI-The-promise-and->

[challenge-of-the-age-of-artificial-intelligence-in-brief-Oct-2018.ashx#:~:text=Embracing%20AI%20promises%20considerable%20benefits,companies%2C%20sectors%2C%20and%20countries.](#)

Miller, Katharine; *Introducing The Foundation Model Transparency Index*, 2023;
<https://hai.stanford.edu/news/introducing-foundation-model-transparency-index>

Ozkaya, Ipek; *The Next Frontier in Software Development: AI-Augmented Software Development Processes*, 2023; <https://ieeexplore.ieee.org/document/10176194>

Ricard, Li, Flint, and Lampert; *Keeping Up With Generative AI, part 1 - The opportunity for Insurers* 2023;
<https://www.oliverwyman.com/our-expertise/insights/2023/aug/how-insurers-can-successfully-use-generative-artificial-intelligence.html>

and

https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2023/aug/Oliver_Wyman_Keeping_up_with_Generative_Artificial_Intelligence_in_insurance.pdf

Ridley, Mark; *How Generative AI will Impact Product Engineering Teams*, 2023;
<https://towardsdatascience.com/how-generative-ai-will-impact-product-engineering-teams-83a5eaa8fc60>

Savkin, Alexis; *ChatGPT AI for Scenario Planning: Generate Plausible Stories*, 2022;
<https://medium.com/@bscdesigner/chatgpt-ai-for-scenario-planning-generate-plausible-stories-34b7730da070>

SimpliLearn; *Top Generative AI Tools: Boost Your Creativity*, 2023;
<https://www.simplilearn.com/tutorials/artificial-intelligence-tutorial/top-generative-ai-tools>

SimpliLearn; *What is Generative AI and How Does it Work?* 2023;
<https://www.simplilearn.com/tutorials/artificial-intelligence-tutorial/what-is-generative-ai>

Spaniol and Rowland; *AI-assisted scenario generation for strategic planning*, 2022;
<https://onlinelibrary.wiley.com/doi/epdf/10.1002/ffo2.148>

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 Ave, Suite 1000
Chicago, IL 60631
www.SOA.org