Actuarial A.I 1.0

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

Any views and ideas expressed in the essays are the author's alone and may not reflect the views and ideas of the Society of Actuaries, the Society of Actuaries Research Institute, Society of Actuaries members, nor the author's employer.

HELLO WORLD!

What would it be like to have a helper that knows all the ins and outs of your model, was able to provide documentation in an instant, and could provide modelling suggestions for your every whim?

The ability to trust and understand a model is as important as it has ever been. Everything seems to be a little more complicated in the actuarial space that it once was: increased volume of regulatory disclosures, more robust capital frameworks, and ever-expanding model sophistication.

The actuary's time is very precious. Adding efficiency by shortening learning curves and facilitating documentation exercises can give more time back to actuaries for analytical work.

At our company, AI was adopted and shared rather quickly. Overnight, 10,000 plus workers became kids with new toys, encouraged to play, invent, and innovate with AI. The future was here.

There were some guardrails. The earliest AI utilities were kept within corporate walls, only allowed to look at our own documents. Still, a flood of ideas came from these early experiments, as well as a new library of prompts and inspiration for many ways to enhance the actuarial experience.

One of the initial projects we took on was an attempt to create an AI "Navigator" that had access to all the documentation and support material behind our actuarial projection software. Would it be the perfect helper? Could it understand actuarial lingo? We set off to find out.

LEARNING CURVES

The software had 30 plus years of development, thousands of pages of help text, and numerous specially focus user guides covering topics like GAAP LDTI, IFRS 17, and Capital Frameworks across the globe. No one person at the company knew everything about the software (although an argument could be made for our Chief Actuary), so this would augment the client support staff in a powerful new way.

We formed a dedicated team to populate a repository of information with existing resources.

We started small, isolating things to one segment of the software before expanding further. Test, revise, repeat.

Some of our early learnings:

- **Partition documentation into related sets of information and direct the tool to look in these specified areas**. This helps reduce false answers ("hallucinations"). For example, if I want to know how the model handles partial withdrawals on an annuity, I don't need to reference the documentation on life reserves.
- **Source material needs to be formatted well, accurate and current**. Some early prompts revealed older processes that weren't the latest best practice or could only provide an inconsistent or incomplete answer.
- **Prompting suggestions are important to develop along with the AI utility.** While the tools are not well versed in actuarial context, providing a pre-prompt that might be considered overly verbose may end up being appropriate. For example, "Consider 'experience study' to be distinct from 'experience adjustment'." This type of clarifying statement can help it distinguish between an assumption setting subject and a reinsurance topic.

Similar lessons were coming from other AI applications being developed and tested, including email thread summarizers, coding agents, and research assistants.

PERFORMANCE REVIEW

Some of the earliest encouraging moments were when the tool compiled documentation from multiple sources in a coherent, well organized conversational summary. Previously, these bits of information were in scattered locations, which caused a drag on time for staff or clients to search for and locate needed reference material. Being able to bring these disparate chunks of information together is a huge win.

Additionally, after initial prototypes were proven to be successful, expanding the tool to incorporate more information was not a difficult extension.

One key piece of functionality is a rating system (a.k.a "training") that would allow a user to give feedback about whether the answer was satisfactory. Different categorial ratings are offered: "not detailed enough" or "incomplete," "incorrect," etc.

The other helpful feature that allows for easier validation is embedded links to source materials via citations.

Some struggles we encountered:

- Alphabet soup – the actuarial world is full of acronyms, and that could become confusing. Is NPR a net premium ratio, a net premium reserve, or something else?
- Blind spots. Because we intentionally restricted the universe of information for the tool, anything that lacked documentation would impede the tool's success. This could lead to negative ratings due to "incomplete" answers, or the tool being unable to provide a response at all.
- Assisting with user-based code presented a challenge because of the required rules and syntax of the underlying software. Continued development is underway to help support this effort. However, we did find that the tool could provide great summaries of what pre-existing code was doing, including clean presentations of input and output variables.

PREPARING FOR PRIME TIME

Before its public release, the tool is being Beta tested by a select group of users. Phased deployment is a very useful technique in software development.

In our outreach, we had nervous enthusiasm. Insurance companies seem to be tilting towards cautious adoption of AI tools. There are understandable concerns about exposing private information or intellectual property to these utilities. Having clear and demonstrable guardrails will be a continued important feature to ensure successful deployment in the industry.

We also needed to provide adequate training, which included some examples of how to use the tool and suggested prompting.

Following customer feedback, we will get back under the hood and make some tweaks and then hopefully release version 1.0.

LOOKING AHEAD

Our hope is that this tool will begin to plant the seeds for that perfect actuarial assistant that can resolve governance woes and facilitate learning. In the meantime, this type of AI engine will deliver a better customer experience by driving higher quality engagement with our support staff. The practice of actuarial science is nuanced and often involves judgement, which is a difficult thing to train into an AI agent. There will continue to be a role for human-to-human actuarial discourse in the foreseeable future, but there will also be plenty of augmented human-to-augmented human discourse as well.

This is just the first step in our actuarial AI journey. Our helper can aide in understanding a concept or feature in the model. There are several areas where we see continued usefulness.

- **Model Documentation**: We know actuaries and programmers *love* documentation. Or maybe they see it as a necessary evil. It is crucial for model governance and certainly having a robot do the heavy lifting here will be a time saver.
- **Model Building**: Beyond actuarial applications, the world of tech is moving to low or no code interfaces. Simply prompt about what you need and let AI do the rest. In the actuarial world, maybe this is a pricing model that can emerge from a product specification. To accomplish this, model providers will need to design/build APIs for agents and write documentation specifically to guide agents.
- **Results Analysis**: From experience studies to period-to-period valuation investigations, frameworks for helping to highlight adverse results will be yet another time saver. This may be especially helpful for large and complex organizations, and also useful when there are many components to reported results (like IFRS-17 and GAAP Roll-forwards).

The future is here. Actuarial and insurance workflows will encounter artificial intelligence all along the way including automation steps, assistance for performing actuarial work, agents to aid in understanding and synthesizing complex information, writers of documentation, and even quantitative analysis. Knowing how to approach these tools, providing solid foundations, iterating frequently, and preparing carefully for launch is critical to ensure success.

*   *   *   *   *

Author Byline: Nathan Worrell, FSA, Green Chen, FSA, and Brandon Lin, FSA, are actuaries at Moody's. They can be reached at Nathan.Worrell@moodys.com, Green.Chen@moodys.com, and Brandon.Lin@moodys.com, respectively.