SOCIETY OF ACTUARIES

Article from

Health Watch

August 2020

A Look Inside the Black Box

By Greg Mottet

It's an exciting time to be a health actuary! The last decade brought with it a wealth of opportunity. And if you blinked, you might have missed some incredible advances in artificial intelligence (AI). Many things that were once impossible are becoming possible. AI has enabled advancements such as the discovery of new planets,[1] the acceleration of drug discovery,[2] the pursuit of treatments for rare diseases,[3] the improved diagnosis and treatment of cancer[4,5] and convenience like never before.[6] It will continue to change the world and will be a defining characteristic of our time. Many of us will tell stories to our grandchildren, and they will be horrified anyone gave us permission to drive to work in a car with a steering wheel.

As this change takes place, new frontiers are emerging. Health actuaries have the opportunity to push health care forward by leveraging AI and machine learning (ML) to improve predictions and provide greater levels of insight. As actuaries learn to blend actuarial judgment with the output of these models, the ability to manage risk will increase, leading to more competitive pricing and better insurance products. For a long time, actuaries have taken advantage of the statistical properties of segments of members to price insurance. But the rich, detailed data underlying these populations has not been fully utilized. Machine learning is the ideal tool to pore over the fine detail behind the data available to actuaries and bring the most critical insights into plain view.

This new frontier also comes with a great deal of responsibility. Used improperly, AI and ML can lead to unintended biases or singling out individuals. For the latter, it is fortunate that regulations exist to protect individuals. For the former, actuaries must be aware of the potential pitfalls and take steps to address them. Care needs to be taken to ensure the model is achieving the intended goal. The actuary must be aware of the potential impacts due to the abundance or sparseness of different types of data and the frequency with which key data is populated over time. Changes in coding patterns, such as the transition from ICD-9 to ICD-10, or the data drift from COVID-19 must be considered. And models must be trained and validated in ways that ensure continuing accuracy on unseen data. The actuary is well suited to address all of these issues. Every actuary has strong domain knowledge, a solid statistical background,

> The most effective actuarial departments and health plans will leverage narrow AI and ML models to improve predictive capability.

technical aptitude and a proven ability to learn new things and tackle complex problems.

THE AI-ENABLED ACTUARY

In its current state, the term "artificial intelligence" is perhaps an overstatement. AI can be categorized along a spectrum of increasing capability starting with artificial narrow intelligence (ANI), progressing to artificial general intelligence (AGI) and ending in artificial super intelligence (ASI). Narrow intelligence has been defined as "the ability to carry out any particular task that is typically considered to require significant intelligence in humans."[7] This is the current state of both AI and ML. Given sufficient data, they can be used to carry out a specific task and can often do incredibly well! They are both well suited to solve problems involving standard tabular data sets, while AI generally has strong advantages in natural language processing and computer vision. But both lack the qualities of general intelligence.

General intelligence is akin to human thinking. Based on one definition, it sounds very much like the work of an actuary: "General Intelligence is the ability to achieve complex goals in complex environments."[8] It has the ability to generalize knowledge from one context and apply it to another. General intelligence is able to understand the impact of regulatory changes. It is able to understand the implications of entering a new market or the presence of a new competitor. It is able to reason through the potential cost of new drugs or medical procedures for which claims data is not yet available.

It will be a long time before AGI is achieved. A recent survey of 352 machine learning and computational neuroscience researchers showed that 50 percent believe that within 45 years "unaided machines [will be able to] accomplish every task better and more cheaply than human workers."[9] But many of those 352 researchers believe AGI will either not be achieved in this century or never be achieved at all.

For the foreseeable future, actuaries will continue to succeed by relying on actuarial judgment and deep domain knowledge. And the most effective actuarial departments and health plans will leverage narrow AI and ML models to improve predictive capability as well as provide continuous retrospective and prospective population insights. Similar to things like provider network discounts, effective underwriting and risk coding, the ability to leverage AI will become a pillar of an effective health plan such that the first health plans to succeed on this front will possess a substantial competitive advantage.

OPENING THE BLACK BOX

A common complaint against machine learning is that the models are a black box. For an actuary, working with a black box is difficult because it limits the ability to assess reasonableness of results, and actuaries need to explain results to business stakeholders and regulators. Complex ML models typically provide significant gains in accuracy compared to linear modeling, but there has always been a trade-off between accuracy and interpretability. Recently developed open-source tools have chipped away at this black box and enabled users to better understand the connection between data inputs and the resulting prediction.

Shapley Additive Explanations (SHAP) is a computationally efficient Python tool that draws on game theory to produce an "interpretable approximation of the original model."[10] It can be used to quantify drivers of individual predictions from complex models, including random forests, gradient-boosted trees and neural networks. The drivers of individual predictions are also additive and can be aggregated across individual predictions to produce population-level insights.

ILLUSTRATION OF SHAP TOOL

For the purpose of illustration, I created a simplified model that predicts member claims for the current calendar year based on data from the prior year and the first month of the current year. The features from prior year data include allowed, paid, paid over allowed, member months and binary diagnosis indicators. The current year features include only paid amounts and allowed amounts for the first month of the year. These features were calculated for a set of members and fed into a gradient-boosting model. The model hyperparameters were then tuned using cross-validation to minimize mean squared error. Next, to avoid displaying the personal data of individuals, two additional rows of data were manually created representing two "pseudo" members. A model like this leaves a lot on the table, but it serves the purpose of illustrating the value of SHAP.

Figure 1 shows the most predictive features in descending order and how various values of those features influence the prediction. The x-axis represents the impact on the per member per month (PMPM) claims prediction. Each point represents a member, and the color scale corresponds to the value of the feature. Areas with higher member density will allow the points to spread out above and below the horizontal line. For example, the point in the upper right represents a member whose prior year allowed was on the high end (bright red), which contributed around $4,000 PMPM to the prediction.

A figure like this can be used to assess the reasonableness of the connection between the data and the predictions. Say we unintentionally added member ID as a feature, and it was highly

Figure 1
SHAP Summary Plot: Explaining Feature Impact on Predictions

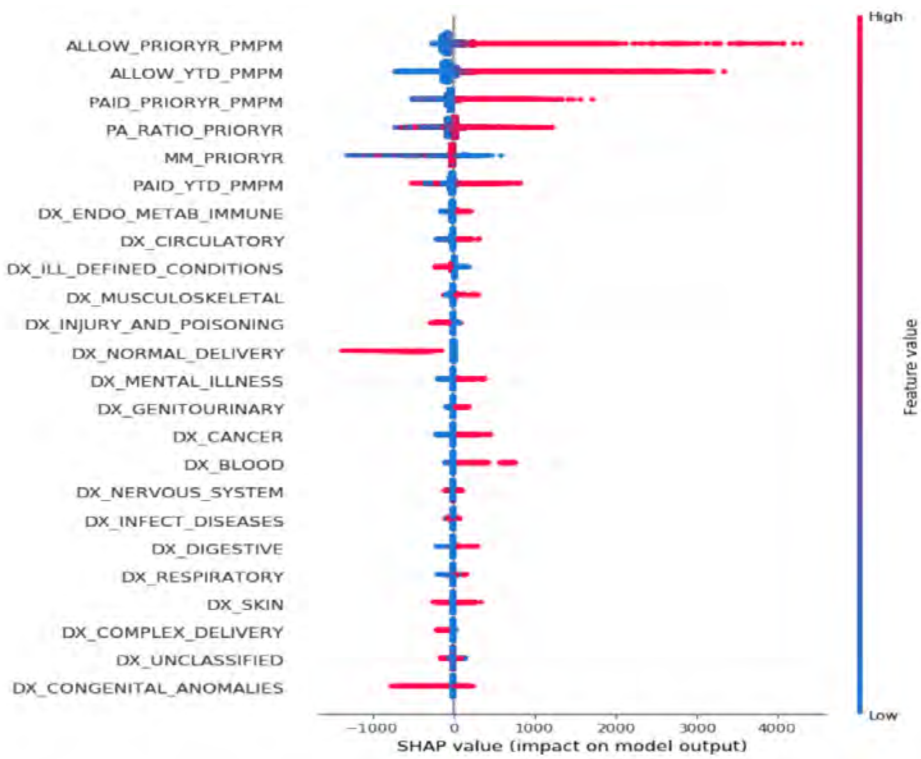

predictive. Looking at this figure would quickly show that the model created a spurious connection between member ID and claim cost. Another insight we can glean from Figure 1 is that the model is lowering cost the year after a delivery occurred. High values of the ALLOW_PRIORYR_PMPM result in higher predictions, but we see that when there was also a delivery in the prior year (DX_NORMAL_DELIVERY = 1), the model learned to reduce the prediction for the subsequent year. What else do you see that would give you reason to believe the model is producing reasonable or unreasonable results? This is the time to apply your judgment.

Figure 2 shows the impact each feature had on the individual predictions for the pseudo members created. The average predicted value across the data set (displayed as the "base value") is $416.90 PMPM. The predicted claims for member A are $493.22 PMPM. A circulatory diagnosis contributed about $80 to the prediction. A diagnosis in the endocrine, metabolic and immune category contributed an additional $50 or so. And a relatively high paid-to-allowed ratio pushed the prediction a little higher. On the other side, there were 12 months of enrollment in the prior year, which lowered the prediction by about $20. Member B had claims of $15 PMPM in the prior year and has yet to have a claim in the current year, which are the top drivers of a low prediction of $98.26 PMPM. Note that the drivers are additive, and provided a link function is not used, results can be directly aggregated to member segments and provide insights into how underlying changes in the population influence overall cost.

Recently developed open-source tools have enabled users to better understand the connection between data inputs and the resulting prediction.

Figure 2
SHAP Force Plot: Explaining Predictions for Individual Members

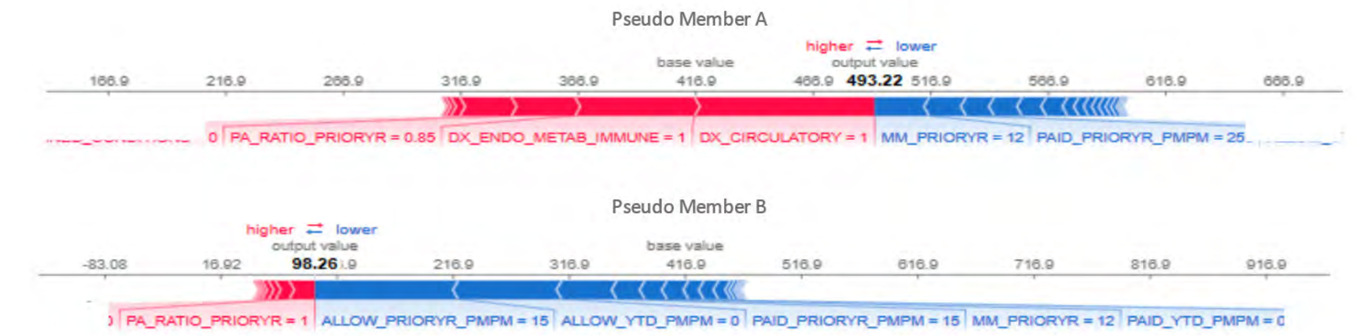

APPLICATIONS FOR HEALTH ACTUARIES

The value gained from applying these methods in health insurance is only beginning to be realized. New applications will continue to emerge, and with enough vision and grit, the health care we have in the future will look much different and better than what we have today. As health actuaries branch out of traditional areas of practice, opportunities to improve health care will expand. But even now, we can start thinking about some of the ways we have done things in the past that can be done better with new tools. In many cases, standard models and actuarial judgment will continue to be the wisest course, while in others, AI and ML models will offer strong advantages. With that in mind, these are some areas where it makes sense to search for innovation in the near term:

- **Financial forecasting and monitoring.** Data from the current and prior years can be used to complete member claims for the current year. Combining these predictions with SHAP output and aggregating to lines of business can provide continuous feedback into drivers of claim cost, strategic insight and improvement of trend estimates.

- **Large group rating and underwriting.** A traditional rating formula trends and adjusts from a group's last year of experience. Credibility is used to control rate fluctuations, but often no attempt is made to control for mean reversion and future rising risk, which is left to the underwriter. The predictive power of diagnosis codes, prescription drugs and procedure patterns can be used to systematically improve both pricing accuracy and long-term rate stability. Further, combining predictions with SHAP output can give underwriters confidence in knowing when to hold the rate or make concessions.

- **Risk coding.** AI and ML models can leverage patterns in data that identify missing or inaccurate diagnosis codes. This information can be used to supplement existing methods for complete and accurate coding of members that would otherwise be hard to identify.

- **Case management.** ML models excel at predicting the high end of risk. They are able to capture more complex interactions than traditional linear models, such as the timing of events in relation to one another.

- **Reserving.** AI and ML models have shown promise in predicting reserves with high accuracy. Interpretable output would also allow actuaries to blend their unique insights from factors not present in data.

THE FUTURE IS BRIGHT

The pace of change and adoption of AI over the past decade has been remarkable. Looking ahead, this trend is very likely to continue. In another decade, I believe the achievements relating to health insurance will be far greater than suggested by the list in this article. As we increase our understanding of AI and ML, new opportunities will be discovered. And as we work across organizations, we will be better equipped to solve some of the toughest challenges in health care. Blending together AI and ML with our domain knowledge and actuarial judgment will cut through barriers that were once difficult to overcome. Some

As health actuaries branch out of traditional areas of practice, opportunities to improve health care will expand.

things that once seemed impossible will become possible. Keep learning and develop a skill set in AI and ML! This trend is here to stay and is certainly an exciting one to be a part of. ■

Greg Mottet, FSA, MAAA, works with Prealize Health developing machine learning models to improve prediction and understanding of health care cost. He can be reached at *gmottet@prealizehealth.com.*

ENDNOTES

1 NASA. Artificial Intelligence, NASA Data Used to Discover Eighth Planet Circling Distant Star. *NASA*, December 14, 2017, *https://www.nasa.gov/press-release/artificial-intelligence-nasa-data-used-to-discover-eighth-planet-circling-distant-star* (accessed July 12, 2020).

2 Brasil, Sandra, et al. 2019. Artificial Intelligence (AI) in Rare Diseases: Is the Future Brighter? *Genes* 10, no. 12:978. DOI:10.3390/genes10120978.

3 Ibid.

4 Walsh, Fergus. AI "Outperforms" Doctors Diagnosing Breast Cancer. *BBC News,* January 2, 2020, *https://www.bbc.com/news/health-50857759* (accessed July 12, 2020).

5 Sokol, Emily. AI Brings Personalized Cancer Treatment to Cleveland Clinic. *Health IT Analytic,* July 11, 2019, *https://healthitanalytics.com/news/ai-brings-personalized-cancer-treatment-to-cleveland-clinic* (accessed July 12, 2020).

6 Karlinsky, Neal. How Artificial Intelligence Helps Amazon Deliver. ***The Amazon Blog,*** June 5, 2019, *https://blog.aboutamazon.com/innovation/how-artificial-intelligence-helps-amazon-deliver* (accessed July 12, 2020).

7 Goertzel, Ben, and Cassio Pennachin, eds. 2007. *Artificial General Intelligence.* Rockville, MD: Springer.

8 Ibid.

9 Grace, Katja, et al. 2018. When Will AI Exceed Human Performance? Evidence From AI Experts. *Journal of Artificial Intelligence Research* 62:729–754, *https://jair.org/index.php/jair/article/download/11222/26431/* (accessed July 20, 2020).

10 Lundberg, Scott M., and Su-In Lee. A Unified Approach to Interpreting Model Predictions. Paper presented at the 31st Conference on Neural Information Processing Systems (NIPS 2017), December 4–9, 2017, Long Beach, Calif., *http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf* (accessed July 12, 2020).