BY BART CUSTERS
AND GARRY MARTES

Deep Learning for actuaries

Recent breakthroughs in artificial intelligence (AI) have been largely driven by neural networks, particularly deep learning frameworks. These techniques have revolutionized fields such as image recognition, natural language processing, and predictive analytics. However, it can be difficult to apply neural networks in actuarial settings. First, neural networks are often considered 'black boxes' because of their complex structure with thousands of parameters. Actuaries usually produce rating structures that clearly show how premiums and risks behave over different variables, as transparency and interpretability are crucial in actuarial work. Neural networks though, only return a single prediction which makes it much more difficult to interpret model results and provide explanations to stakeholders. Second, most actuarial datasets have a tabular structure, consisting of structured numerical and categorical variables. On tabular data, traditional models such as Generalized Linear Models (GLM) or tree-based models (e.g. Random Forest, GBM) usually perform better compared to neural networks.

B. Custers MSc AAG (left) is Senior Pricing Actuary at Allianz Direct and Student MSc Artificial Intelligence.

G. Martes MSc AAG is Pricing Actuary at Allianz Direct.

This article has been written in a personal capacity.

Recognizing these challenges, deep learning frameworks suited for actuarial problems and tabular data have been developed in the academic community. In this article, we discuss three of these frameworks: the Combined Actuarial Neural Network (CANN), LocalGLMNet and TabNet. We briefly explain the workings of these models as well as the pros and cons of each method.

A BRIEF INTRO TO NEURAL NETWORKS

A neural network (see figure) is a model inspired by the structure of the human brain. It consists of multiple layers of interconnected nodes (neurons) that process the data and aim to learn the relationships in the data. The three main components of a neural network are:

- *Input layer*: This is the layer that processes the input data. Each neuron in this layer represents a feature in the dataset (e.g., policyholder age, claim history, insured amount).
- *Hidden layers*: The hidden layers perform complex transformations on the data. Each neuron in a hidden layer takes inputs from the previous layer, applies a transformation and weighting, and then passes the result through an activation function.
- *Output layer*: The output layer produces the final prediction.

A typical neural network architecture

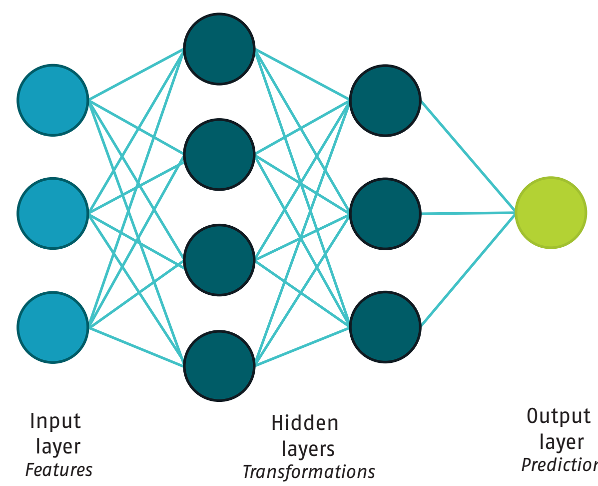

The neural network then learns by adjusting the weights of connections between the neurons to optimize the loss function and predictive accuracy.

THE COMBINED ACTUARIAL NEURAL NETWORK (CANN)

The CANN model, as proposed by Wüthrich & Merz (2019) combines a GLM with a neural network. By using the GLM predictions, the neural network part aims to capture additional patterns that were not captured by the GLM. In the output layer, the final CANN prediction is then obtained by combining the GLM prediction (that arrives via a 'skip connection') with the neural network prediction. In this way, CANN retains the interpretability of GLMs while improving predictive performance.

LOCALGLMNET

LocalGLMNet, developed by Richman & Wüthrich (2023), is another hybrid model designed for actuarial applications. Instead of producing a single set of fixed coefficients like a traditional GLM, it learns for each data record a distribution of local coefficients. For this, the model uses attention weights, which are used in a neural network to determine the importance of different parts of the input data. By multiplying the attention weights with the input data, the structure of a GLM is maintained while allowing for more complex interactions. Furthermore, the attention weights can be used to inspect variable importance and provide local explanations that can help actuaries to explain results.

TABNET

After CANN and LocalGLMNet, we can increase performance and complexity by using the TabNet model. Neural networks are known to underperform on structured data, and TabNet is a deep learning model designed specifically for tabular data. TabNet, as designed by Arik & Pfister (2021), selects the most relevant features through multiple decision steps, using a so-called attention-based feature selection. Unlike traditional neural networks that process all features simultaneously, TabNet only selects important features at each step. This is efficient and reduces the probability of overfitting. In addition, TabNet provides 'masks' that indicate at each decision step which features were found to be important and contributed to the prediction.

EVALUATION

To evaluate these models, we applied them to an open source dataset for car insurance frequency and compared performance to a GLM. The results indicate the huge potential of these methodologies. For readers who want to explore the simplified implementation in detail we provided our code and results on github.com/bart-custers/DL_for_Actuaries.

For actuaries that aim to implement these methods, performance might not be the only criterium to be considered. In the table below we aim to evaluate the deep learning techniques along four criteria: interpretability, predictive power, ease of implementation and computational efficiency. The comparison shows a classic trade-off: simpler models offer better interpretability and straightforward implementation, while more sophisticated models deliver higher performance at the cost of increased complexity and implementation challenges.

	CANN	LocalGLMNet	TabNet
Interpretability	High - Retains GLM structure	Moderate - Attention-based weighting for GLM coefficients	Low - Produces decision masks but more difficult to interpret
Predictive power	Moderate - Able to improve GLM while staying structured	High - Able to improve GLM	High - Outperforms other models in many use cases
Ease of implementation	High - Easy to implement for actuaries	Moderate - Requires hyperparameter tuning	Low - Consists of many parameters that need careful tuning
Computational efficiency	High - Simple structure, no computational issues	Moderate - Heavier in computation due to attention mechanism	Low - Can be computationally very expensive

CONCLUSION

Deep learning with neural networks has transformed many industries, but adoption in the actuarial field could be limited due to concerns about interpretability and suitability for tabular data. Models such as CANN, LocalGLMNet and TabNet aim to overcome these concerns by either combining the strengths of GLMs and neural networks or by specifically focusing on deep learning for tabular data. Our results demonstrate that these models have great potential to improve predictive performance while maintaining interpretability. Therefore, these techniques show a promising direction for actuaries that aim to leverage deep learning in their work. ■

References

Sercan Arik and Tomas Pfister. 2021. TabNet: Attentive Interpretable Tabular Learning. *In The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*. 6679-6687.

Ronald Richman and Mario V. Wüthrich. 2023. LocalGLMnet: interpretable deep learning for tabular data. **Scandinavian Actuarial Journal** 2023(1), 71-95.

Mario V. Wüthrich and Michael Merz. 2019. EDITORIAL: YES, WE CANN! **ASTIN Bulletin** 49(1), 1-3.