


data symphony

Creating Business Value, Driven by Data Intelligence

CASE STUDY

Unlocking Actuarial Model Precision with
eXplainable Neural Networks (XXNs)



www.datasymphony.com



South Africa | Australia



Case Study

Unlocking Actuarial Model Precision with eXplainable Neural Networks (XXNs)



Discover how XXNs improve the accuracy of mortality predictions while maintaining the transparency needed for regulatory compliance and business decision-making in the insurance industry.



Overview

In the rapidly evolving world of actuarial science, data-driven models are at the heart of critical business decisions. Traditionally, Generalized Linear Models (GLMs) have been the benchmark for mortality rate predictions in the life insurance industry. However, with the advent of deep learning and advanced computational techniques, more powerful and flexible models have emerged. This case study explores the application of eXplainable Neural Networks (XXNs) to actuarial mortality data, showcasing how these models can improve prediction accuracy while maintaining the critical transparency and interpretability required for regulatory compliance and informed business decision-making.



Problem Statement

The insurance industry operates within strict regulatory frameworks, including Solvency II and SAM (South African Market), which mandate model transparency to ensure the robustness and fairness of pricing, risk assessment, and capital requirements. The challenge is that traditional GLMs, while interpretable, often struggle to capture the complex, non-linear relationships present in mortality data. At the same time, while neural networks offer enhanced predictive power, they are typically seen as "black-box" models that lack interpretability.

The goal of this case study is to bridge this gap by applying XXNs, a deep learning model that maintains the predictive power of neural networks while incorporating explainability. This ensures that actuaries and business stakeholders can understand, trust, and act upon the model's insights.





Challenge

Interpretability and Regulatory Compliance: The insurance industry operates under a rigorous regulatory landscape, with frameworks such as Solvency II, IFRS 17, and Own Risk and Solvency Assessment (ORSA) demanding transparency in modeling practices. These regulations require that actuarial models be fully interpretable, enabling actuaries to explain how each factor influences the outcomes. This is particularly vital when models are used for risk-adjusted capital calculations, stress testing, and underwriting. Non-compliance with these regulatory demands can lead to significant repercussions, such as financial penalties, regulatory sanctions, or even product withdrawal from the market. Therefore, ensuring that advanced models like XXNs meet these interpretability standards is critical for regulatory adherence, while simultaneously providing the necessary detail to support auditability and traceability within actuarial processes.

Complexity of Non-linear Relationships in Mortality Data: Mortality data presents highly complex, non-linear relationships across various dimensions. Variables such as age, gender, health status, and behavioral factors like smoking or alcohol consumption interact in intricate ways. For example, a smoker's mortality risk may change dramatically based on their profession or geographic location due to environmental factors, but these non-linear interactions are not well captured by traditional models like GLMs. Furthermore, the impact of certain factors may not follow a linear path; for instance, health interventions may reduce risk for certain age groups but exacerbate risks for others. These heterogeneous effects pose a substantial modeling challenge, where advanced techniques like XXNs are needed to recognize and respond to these dynamic, non-linear relationships without overfitting or misinterpreting the data.






Challenge

Balancing Predictive Accuracy with Explainability: While neural networks, and specifically XXNs, provide the ability to significantly boost predictive accuracy, they introduce a level of complexity that can hinder interpretability. The "black-box" nature of most neural networks makes it difficult for actuaries to pinpoint why the model reached a certain decision. In the insurance sector, where actuaries must justify pricing, risk assessments, and reserving calculations, having a highly accurate but opaque model can be detrimental. The challenge is to maintain the granular insights and precision offered by XXNs while implementing methods to provide clear, intuitive explanations. This balance is crucial, not only for actuarial approval but also for gaining trust from non-technical stakeholders such as C-level executives, product managers, and regulators.

Risk of Overfitting and Model Robustness: Neural networks, while highly flexible, also carry the risk of overfitting, especially when applied to highly granular mortality data. Overfitting occurs when a model performs exceptionally well on training data but fails to generalize effectively to unseen data, leading to poor performance in real-world applications. This is particularly concerning in the insurance industry, where robustness and stability in predictive models are critical for managing long-term liabilities. Actuaries need to ensure that the model is calibrated properly, avoiding overfitting while maintaining the flexibility to adapt to changing data patterns over time, such as shifts in mortality trends due to pandemics, medical advancements, or economic changes.





Solution Approach

Dataset Overview

The mortality dataset used in this study, sourced from Kaggle, includes anonymized policyholder demographic data and risk factors that influence mortality rates. The dataset covers variables such as age, gender, profession, region, and smoking status—all key indicators of life expectancy and risk in insurance.

The traditional approach uses Generalized Linear Models (GLMs) to predict mortality rates, but these models struggle with the non-linear relationships inherent in such datasets. Our solution involves applying eXplainable Neural Networks (XXNs), which retain the advanced learning capabilities of neural networks while allowing for interpretability through various tools.

Tools and Libraries

To implement XXNs, a combination of Python libraries was used:

- TensorFlow and Keras: For constructing and training the neural network.
- Pandas and NumPy: For data manipulation, cleaning, and preprocessing.
- Matplotlib and Seaborn: For data visualization, helping to understand the structure of the dataset.
- SHAP (SHapley Additive exPlanations): To interpret the model's outputs, ensuring transparency.
- Scikit-learn: For preprocessing and evaluating models, including splitting data into training and test sets.
- statsmodels: For fitting GLMs, providing a benchmark for comparison.






Solution Approach

1. Data Preprocessing and Model Development

- **Data Cleaning:** The initial step in preparing the mortality dataset involved thorough data cleaning and preprocessing to ensure its suitability for modeling. This process began with handling missing values by removing any rows that contained significant gaps in information, thereby enhancing the dataset's integrity.
- **Categorical Variable Transformation:** Categorical variables, such as policyholder profession and region, were transformed into a numerical format using one-hot encoding. This technique allows the model to interpret categorical data effectively by creating binary columns for each category.
- **Normalization of Continuous Variables:** Continuous variables, such as age and income, were then normalized to ensure that they were on a similar scale. This step is crucial for improving the convergence speed of the models during training, allowing the algorithms to learn more efficiently.
- **Data Splitting:** After preprocessing, the dataset was divided into training, validation and testing sets to evaluate model performance accurately. This division ensures that the model can be trained on one subset of the data, while tuning hyperparameters on another and finally be tested on another unseen set of data, helping to prevent overfitting and information leakage.





Solution Approach

2. Baseline GLM Model

Model Fitting: To establish a benchmark for evaluating the performance of the XXN model, a Generalized Linear Model (GLM) was fitted using the preprocessed dataset. The GLM serves as a traditional method in actuarial science, providing a straightforward approach to mortality prediction. This baseline model enabled a comparison with the more complex XXN to assess improvements in predictive accuracy.

3. XXN Model Development

Model Construction: For the primary model, an eXplainable Neural Network (XXN) was constructed to harness the power of deep learning while ensuring interpretability. This model featured multiple layers designed to capture complex, non-linear relationships within the data.

Model Training: The XXN was trained using the training dataset, where it learned to identify patterns and correlations between various input features and the predicted mortality rates. The training process involved adjusting the model parameters iteratively to minimize prediction errors, thus enhancing its accuracy and robustness.

Model Validation: Throughout the training, the model was validated using a portion of the data set aside for this purpose. This approach ensured that the XXN could generalize its learning to unseen data, ultimately resulting in a model that not only provided better accuracy but also remained interpretable for actuarial needs.






Model Interpretation Techniques

Interpreting the predictions of complex models like eXplainable Neural Networks (XXNs) is critical for ensuring transparency, trust, and regulatory compliance. To make the model's decision-making process clear and understandable, a combination of advanced interpretation techniques was employed. These methods help both actuaries and non-technical stakeholders grasp how the model arrives at its predictions.

SHAP (SHapley Additive exPlanations)

SHAP assigns a numerical value to each feature, representing its contribution to the model's prediction for a specific outcome. These values quantify the impact that individual features have on the predicted mortality rates, allowing for a clear understanding of the model's logic.

- **Application:** SHAP was used to identify key drivers of mortality predictions, such as age, smoking status, and profession. For example, SHAP can demonstrate that being a smoker increases mortality risk more significantly than being in a particular profession.
- **Benefit:** SHAP not only enhances transparency but also helps actuaries explain complex model behaviors in a simple, quantifiable way. This is especially important when presenting to regulatory bodies or business leaders, as it provides confidence in the model's results and justifications for decisions related to pricing and risk management.





Model Interpretation Techniques

Partial Dependence Plots (PDPs)

PDPs help visualize how individual features affect the model's predictions. They show the average predicted outcome as a specific feature changes, while keeping all other features constant. This makes it easier to see non-linear effects in the data.

- **Application:** PDPs can illustrate how age influences mortality rates. A plot might show that mortality rates increase sharply after age 60, highlighting risk factors that traditional linear models like GLMs might fail to capture accurately.
- **Benefit:** PDPs provide an intuitive way to explain the effect of specific features, making it easier to understand and communicate non-linear relationships. This is valuable for actuaries who need to translate technical insights into practical, business-oriented recommendations.

LIME (Local Interpretable Model-Agnostic Explanations)

LIME provides local explanations for individual predictions by approximating the neural network with simpler models in the vicinity of a specific instance. This helps to explain why a particular prediction was made for a given data point, without needing to interpret the entire model globally.

- **Application:** LIME was used to explain individual predictions, such as why a particular policyholder's mortality risk was classified as higher or lower. It breaks down the specific features that most influenced that decision, such as the policyholder's age, health status, or smoking history.
- **Benefit:** LIME is particularly useful for understanding edge cases or unusual predictions that may deviate from expected trends. This allows actuaries to drill down into specific instances, providing actionable insights and justification.





Outcomes

- **Improved Prediction Accuracy** - The XXN model achieved a 16% improvement in prediction accuracy compared to the baseline GLM model. Accuracy was evaluated using industry-standard metrics, including Mean Squared Error (MSE) and R-squared, demonstrating the XXN's superior ability to capture non-linear relationships between mortality risk factors.
- **Enhanced Non-Linear Modeling** - The XXN model was particularly effective in capturing non-linear relationships within the data, such as the impact of age, smoking status, and profession on mortality rates. For example, it was able to identify the mortality hump in early adulthood and the sharp divergence in mortality risk between smokers and non-smokers—patterns that are difficult for linear models like GLMs to capture accurately.
- **Interpretability with SHAP and LIME** - SHAP provided global insights, showing how each feature (e.g., age, smoking) influenced overall predictions, while LIME provided local, instance-specific explanations, allowing actuaries to justify individual mortality risk assessments. The use of SHAP and LIME allowed for a detailed explanation of how individual features contributed to the model's predictions.
- **Risk Management and Increased Business Value** - The XXN model's ability to predict mortality risk more accurately supports more precise risk management and capital allocation strategies, ensuring that insurance companies can better assess long-term liabilities and capital reserves. The combined use of XXNs and interpretability tools provided a competitive edge by improving the accuracy and transparency of mortality predictions, driving better business outcomes through more accurate pricing and risk mitigation.



data symphony

Creating Business Value, Driven by Data Intelligence



GET IN TOUCH
ask@datasymphony.com



www.datasymphony.com

