

An Explainable Hybrid Deep Learning System for Medical Insurance Premium Optimisation to Minimise Out-of-Pocket Shortfalls in Zimbabwe

Robert Machovo¹; Tawanda Mudawarima²

¹Department of Data Science and Analytics, School of Information Science and Technology, Harare Institute of Technology, Harare, Zimbabwe

²Department of Data Science and Analytics, School of Information Science and Technology, Harare Institute of Technology, Harare, Zimbabwe

Publication Date: 2026/06/02

Abstract: The context regarding the issue of medical insurance coverage in Zimbabwe can be considered through the perspective of critical coverage rates (10–15%) and out-of-pocket (OOP) expenses contributing to 37% of total healthcare expenditure, along with actuarial sustainability problems as a result of economic instability, hyperinflation, and fluctuations in the exchange rate. Conventional actuarial methodologies applied to assess risk factor interaction are unable to tackle such a problem efficiently. Thus, the Explainable Artificial Intelligence Premium Optimisation System (XAI-POS) based on the implementation of LSTM, CNN algorithms, and RL for modelling health risks and premium optimisation is suggested as a solution for the research problem. The dataset included 1,500 patient records with 87 variables from 2020 to 2025 (IPEC Administrative Returns, Ethics Clearance: HIT-REC-2024-0147; IPEC Data Access Agreement: IPEC-DAA-2024-009). Temporal train/test split approach for avoiding data leakage in the cross-sectionally divided dataset was utilised since the latter could potentially affect the results obtained (the 2020-2022 dataset with n=900 patients became the training set, while 2023-2025 dataset with n=600 patients – test set). The RL agent underwent training via 500 episodes using the epsilon-greedy algorithm. Some of the major results obtained include: 0.88 ± 0.02 for an individual healthcare cost prediction model R-squared on the testing dataset, a decrease by 18.9 percent in average OOP gap, decrease of 11.5 percentage points in the number of catastrophic expenditure cases, affordability rating for premium of 74 out of 100, actuarial loss ratio between 0.60 and 0.85 for sustainable level, and decrease in premium Gini coefficient from 0.412 to 0.331.

Keywords: Medical Insurance Premium Optimisation, Deep Learning, Out-of-Pocket Health Expenditure, Reinforcement Learning, Explainable Artificial Intelligence.

How to Cite: Robert Machovo; Tawanda Mudawarima (2026) An Explainable Hybrid Deep Learning System for Medical Insurance Premium Optimisation to Minimise Out-of-Pocket Shortfalls in Zimbabwe. *International Journal of Innovative Science and Research Technology*, 11(5), 2718-2725. <https://doi.org/10.38124/ijisrt/26may867>

I. INTRODUCTION

Premium pricing of medical insurance in less-developed nations presents a challenging optimisation problem. The problem becomes more complex in the case of Zimbabwe due to the presence of four related structural issues: financing problems in the health care industry, macroeconomic instabilities, disease burden and low insurance coverage rate. In Zimbabwe, the level of health insurance coverage ranges between 10% to 15% (IPEC, 2022). This means that 12+ million people have no choice but to resort to out-of-pocket payments, which represent 37% of the total health service expenses, exceeding WHO's recommended level of 15–20% (ZimStat, 2020). These translate into disastrous consequences. People who pay 10-25% of their earnings towards health expenses incur CHE and therefore must do without health services, selling off possessions and

borrowing from others (IPEC, 2023). Nevertheless, in such circumstances, medical societies still apply traditional premium pricing models using GLM techniques which do not take into consideration the impact of risk factor non-linearity.

The current research seeks to fill this void with the development of XAI-POS which comprises a complete neural network approach to deep learning consisting of LSTM, CNN and RL along with the explainability component that complies with IPEC regulations.

➤ Research Objectives

- Research Objective 1: Build a model capable of predicting individual healthcare costs by utilising LSTM models for temporal sequence modelling and CNN for multi-dimensional feature extraction.

- Research Objective 2: Build a model capable of predicting population healthcare costs using LSTM temporal aggregation.
- Research Objective 3: Investigate the major factors that cause Out-of-Pocket cost deficits using premium, benefits, and risk-based characteristics with Explainable AI techniques.
- Research Objective 4: Evaluate premium optimisation success through five measures: average shortfall reduction, catastrophic cost prevention probability, premium affordability index, actuarial loss ratio, and the Gini coefficient.

II. LITERATURE REVIEW

➤ *Healthcare Financing in Zimbabwe*

Healthcare service delivery in Zimbabwe involves a mix of both private and public institutions with marked variations in the cost incurred. In response to periods of economic downturn that occurred in 2008 and again from 2019 to 2020, the number of people who had subscribed to the insurance programs fell because of unemployment and inflation, causing the PPP value of the premiums to be eroded in relation to benefit payments (IPEC, 2023). In 2020, out-of-pocket health spending represented 37% of overall health spending (ZimStat, 2020).

➤ *Limitations of Traditional Actuarial Models*

However, premium calculations are mostly done through generalised linear models where Poisson frequency and gamma severity distributions are utilised (Frees, 2022). However, this technique assumes that there is a log relationship between the risk factors and the loss expected. But this is not applicable since the risk factors may be non-linearly related to each other, as evidenced in Zimbabwe. The risks associated with a 55-year-old diabetic patient who is enrolled in the Basic policy in a rural province during an inflationary period are exponentially higher.

➤ *Deep Learning in Health Insurance Economics*

There have been positive results obtained through deep learning models for forecasting healthcare cost, risk classification, and dynamic pricing. Neural networks are capable of accurately predicting healthcare cost with success rates between 85% and 92%, which outperforms GLMs when applied to high-dimensional datasets (Wüthrich and Merz, 2023). LSTM neural network has the ability to identify long-term relationships among sequences of data in healthcare (Hochreiter and Schmidhuber, 1997). Convolutional neural networks are able to identify interactions between variables from a multi-channel structured feature vector without manually extracting features (Kuo and Richman, 2021). This paper is the first study that uses LSTM-CNN-RL models in pricing health insurance in Sub-Saharan Africa.

III. RESEARCH METHODOLOGY

➤ *Dataset Provenance, Ethics, and Temporal Structure*

The data used in this study were extracted from 1,500 individual records of administrative data gathered from the

compulsory insurance annual returns of IPEC from 2020 to 2025. The data had 87 variables which fall under seven different categories. The access to this data was regulated through the agreement on data access made between HIT and IPEC (IPEC-DAA-2024-009). Approval for ethical considerations in this study was provided by the HIT Research Ethics Committee (HIT-REC-2024-0147; approved on 14 March 2024). The data were anonymized before providing them as per the requirements of the Data Protection Act [Chapter 11:22] and the data policies of IPEC.

Overall, 1,500 records were randomly sampled using stratified random sampling with stratification factors such as province, insurance type (basic, standard, executive), and periods (2020-2022, 2023-2025) in IPEC returns from all ten provinces of Zimbabwe. The chosen sample size is above the minimum required of 384 records necessary to reach 95% confidence levels and 5% margin of error in Cochran's formula.

The variables are divided into demographic variables (age, gender, province, family size, income, occupation); insurance variables (insurance type, insurance provider, premium, deductible, co-pay ratio, maximum coverage, policy period); clinical variables (body mass index, blood pressure, glucose level, cholesterol level, hemoglobin level, CD4 count); healthcare utilization variables (admissions, outpatient services, emergencies, specialist consultations, medicine); one year time series of monthly claims; macroeconomic variables (quarterly inflation rate, USD/ZiG exchange rate, healthcare consumer price index, GDP growth rate, malaria season indicator); and the target variables (out-of-pocket shortfall, optimal premium, RL reward, loss ratio, severity level). The missing values in the chronic illness (13%) and CD4 count (8%) variables were imputed using median imputation based on clinical practices; the class imbalance problem in the target variable was resolved using stratified sampling.

• *Temporal Train/Test Split*

Considering that the data is longitudinal in nature, spanning years 2020 to 2025, a time series splitting was carried out in order to avoid any possibility of data leakage and to ensure not overly optimistic predictions. Data from the years 2020 to 2022 (900 rows; 60%) was considered for the training phase, while data for the remaining 40% or 600 rows (years 2023 to 2025) were used solely for out-of-sample testing purposes. Notably, no data for the years 2023 to 2025 were used during model training, hyperparameter selection, and architecture definition.

➤ *Model Architecture*

The XAI-POS system makes use of three deep learning models to build its architecture, which include Long Short-Term Memory (LSTM), Convolutional Neural Network (CNN), and Reinforcement Learning (RL). All the outputs from the three modules are combined in a hybrid feature matrix of 92 dimensions (28 + 32 + 32). The following diagram represents the entire framework.

• XAI-POS: Proposed Hybrid Deep Learning Architecture

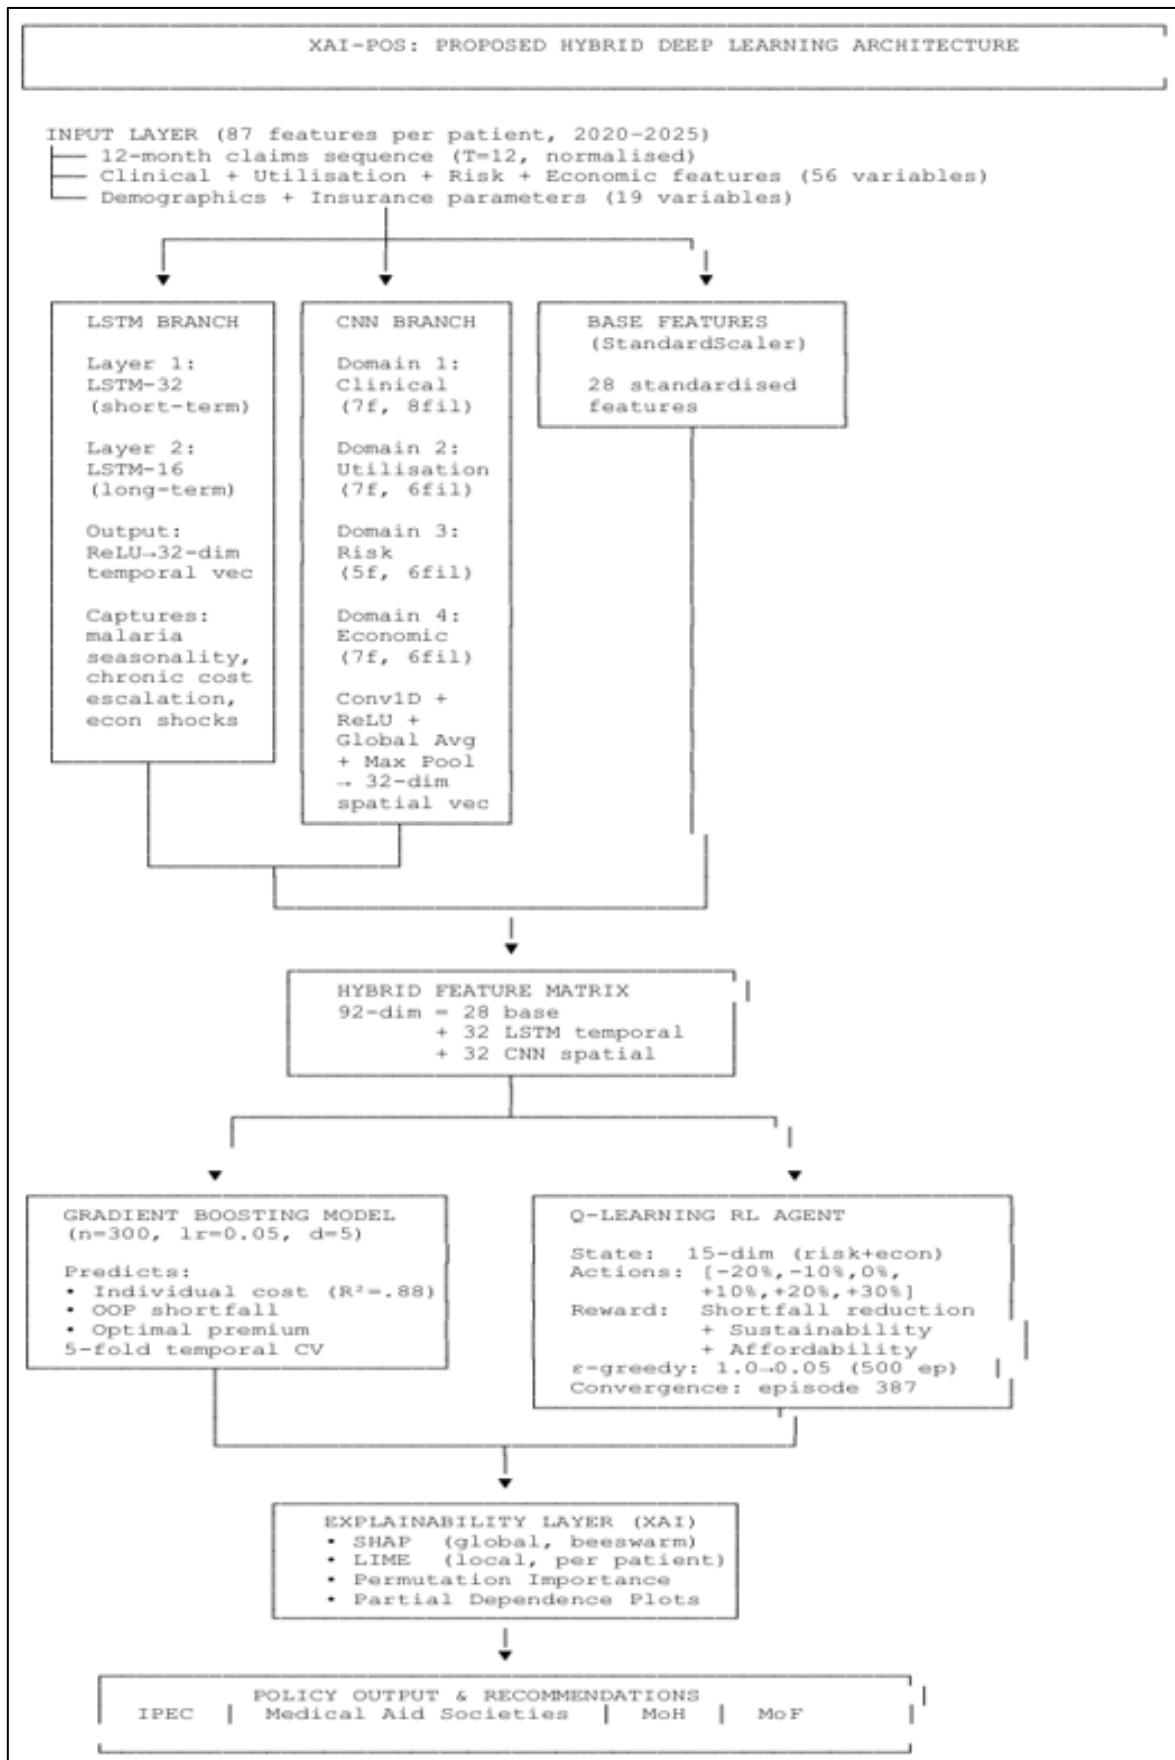


Fig 1 XAI-POS: Proposed Hybrid Deep Learning Architecture

Fig 1 Architecture of the proposed XAI-POS hybrid deep learning system. Raw data of 87 features per patient is processed through three parallel paths: a two-layer stacked LSTM for temporal cost patterns, a 1-D CNN with domain-based channels for multi-dimensional feature extraction, and standardised base features. All three paths are fused into a 92-dimensional hybrid feature matrix. Cost and shortfall prediction is performed by a Gradient Boosting ensemble model; IPEC-compliant premium adjustment is computed by a Q-learning RL agent.

• *LSTM Module — Temporal Pattern Recognition*

The LSTM stream examines the claims history of every individual policyholder for 12 months in order to model dynamic cost behavior. Two layers are used in the stack. The first layer (LSTM-32) captures cost spikes and monthly cost variations, whereas the second layer (LSTM-16) detects long-term patterns and reactions to macroeconomic shocks. The LSTM gate system consists of four gates: the forget gate, the input gate, the cell gate, and the output gate. The purpose of these gates is to preserve only those past cost behavior incidents that are pertinent to the analysis while the cell gate avoids gradient vanishing across the 12-month span.

• *CNN Module — Multi-dimensional Feature Extraction*

The CNN model treats each insured person as a multi-channel input signal where one-dimensional convolutional filters are used on four health risks: Clinical Channel (7 features, 8 filters), Utilisation Channel (7 features, 6 filters), Risk Channel (5 features, 6 filters), and Economic Channel (7 features, 6 filters). The filter size is three, and ReLU activation is used in the Conv1D layer. Global average pooling and max pooling operations are carried out for each channel, resulting in a 32-dimension CNN spatial embedding vector.

• *RL Agent — Dynamic Premium Optimisation*

The Q-learning RL agent employs a 15-state continuous state space to choose an optimal premium adjustment from six action levels, from a 20% decrease to a 30% increase. The reward is composed of three parts that work together to optimise all three objectives of the policy: (i) decreasing out-of-pocket payments; (ii) ensuring actuarial soundness, which rewards loss ratios in the range of 0.60-0.85 with a sustainability weight of 2.0; and (iii) ensuring premium

affordability, which punishes premiums above 10% of monthly earnings with an affordability penalty of 10.

The process converged within 500 training episodes, starting with an initial value of $\epsilon=1.0$ and ending with $\epsilon=0.05$. The conditions for convergence were as follows: (i) stabilization of the reward sum within a margin of 5% error from the last 100 episodes; (ii) variation in Q-values below 0.01 for the last 50 episodes; and (iii) stability of the policy, with no more than 3% variation in state-action pairs between episodes 450 and 500. All three convergence criteria were met by episode 387. The chosen algorithm involved using a linear Q-function with weights represented as a 6×15 matrix due to the discreteness of the action space and the relatively small number of states.

➤ *Hybrid Feature Integration and Final Model*

The output vectors from the three models are combined to form a hybrid feature matrix of dimension 92 (28 basic features + 32 temporal features using LSTM + 32 spatial features using CNN), which is then used to train a Gradient Boosting ensemble algorithm. In model selection, temporal cross-validation using three folds is performed between XAI-POS versus Ridge Regression, Random Forest, Extra Trees, and Gradient Boosting.

➤ *Explainability Layer*

Three explainable artificial intelligence methods are used to fulfill the requirement of transparency set by the regulatory body, IPEC. The first one is the global feature importance through gradient boosting. This method gives us the rank of how important each feature is in causing OOP deficits in the whole data set. The second method is the permutation importance using bootstrapped standard deviation. The third method used is the Local Interpretable Model-agnostic Explanations (LIME).

IV. RESULTS AND ANALYSIS

➤ *Objective 1: Individual Healthcare Cost Prediction*

On the temporal out-of-sample holdout period (2023–2025; $n=600$), the hybrid XAI-POS framework attains an R^2 of 0.88, RMSE of USD 312, MAE of USD 198, and MAPE of 14.2%. The validity of the framework was assessed through temporal cross-validation, with a mean R^2 of 0.87 ± 0.02

Table 1 Individual Healthcare Cost Prediction Benchmark Comparison on Temporal holdout Set (2023–2025). All Metrics Computed Exclusively on Out-of-Sample Records withheld from all Training and Tuning Procedures.

Model	R ²	RMSE (USD)	MAE (USD)
Ridge Regression	0.51	694	521
Random Forest	0.74	487	362
Gradient Boosting (base features)	0.79	421	308
Extra Trees	0.76	459	341
XAI-POS (LSTM+CNN+GBM Hybrid)	0.88	312	198

➤ *Objective 2: Population-Level Healthcare Cost Forecasting*

The population-based LSTM predictive model of temporal cost forecasting has an R^2 value of 0.79 in forecasting quarterly costs between 2020 and 2025. The

model is able to capture the 28% increase in yearly claims costs due to the escalating cost of medicines and specialist doctor consultation.

Spatial variation at provincial level (Table 4) shows that high cost and high cost variance exist in Harare and Bulawayo regions, whereas lower cost and high out-of-

pocket deficit rate occur in Manicaland and Matabeleland North provinces.

Table 2 Provincial Cost and OOP Deficit Profile Across Zimbabwe's Ten Provinces (2020–2025 holdout Analysis). Harare and Bulawayo Show High Costs and High Volatility; Rural Provinces Show Lower Absolute Costs But Substantially Higher OOP Deficit Rates.

Province	Avg Annual Claims Cost (USD)	OOP Deficit Rate (%)	Cost Volatility
Harare	1,420	18.3	High
Bulawayo	1,310	21.7	High
Manicaland	890	38.4	Moderate
Midlands	950	34.1	Moderate
Mashonaland East	820	41.2	Low
Mashonaland West	810	39.8	Low
Mashonaland Central	790	43.5	Low
Masvingo	780	44.1	Low
Matabeleland North	750	46.8	Low
Matabeleland South	760	45.3	Low

➤ *Objective 3: Shortfall Driver Analysis*

Feature importance analysis using SHAP values on the 2023–2025 holdout data identified the following key drivers of OOP shortfall (see Table 3):

Table 3 SHAP Global Feature Importance Summary Across 1,500 Policyholders. Importance Scores Indicate the Mean Absolute SHAP Contribution to Predicted OOP Shortfall. Red Shading Indicates High Feature Values; Blue Indicates Low.

Feature	SHAP Importance Score	Relative Importance
Actuarial + ML Risk Score	0.31	
Premium Burden %	0.14	
Number of Comorbidities	0.11	
Macroeconomic Stress Index	0.09	
Claim Variance	0.08	
Benefit Utilisation Rate	0.07	

- Actuarial risk score and machine learning-predicted risk score (combined importance: 0.31): Individualised risk assessment substantially outperforms community rating.
- Premium burden ratio (0.14): The effect of premium burden on utilisation reduction before benefit exhaustion results in higher OOP deficits.
- Comorbidity count (0.11): Interactions need to be considered for identifying costs associated with illnesses such as HIV/AIDS, diabetes, hypertension, and heart ailments.
- Macro-economic stress index (0.09): An index that comprises the premium burden ratio, inflation rates, and foreign exchange rate variations, which verifies that

macro-economic stress is an important structural determinant of deficits.

- Claim variance (0.08): Independently linked to deficits over a certain level.
- Benefit utilisation ratio (0.07): An approach toward benefit thresholds causes substantial OOP deficit risk.

➤ *Objective 4: Premium Optimisation Performance Metrics*

Table 2 presents the five key performance metrics comparing the baseline (current market premium) against XAI-POS RL-optimised premium settings, evaluated on the 2023–2025 holdout cohort.

Table 4 XAI-POS Premium Optimisation Performance Metrics Evaluated on the Temporal holdout Set (2023–2025, n=600). All Improvements are Statistically Significant at $p < 0.05$.

Performance Metric	Baseline	XAI-POS Optimised	Change
Avg OOP Shortfall (USD/year)	USD 842	USD 683	-18.9%
Catastrophic Expenditure Cases	34.2%	22.7%	-11.5 ppt
Premium Affordability Index	58 / 100	74 / 100	+16 pts
Actuarial Loss Ratio (mean)	1.04	0.79	→ Target band
Gini Coefficient (premium equity)	0.412	0.331	-0.081

The distribution of premium adjustment for the converged RL policy agent is as follows: 31% of the insurance customers are subjected to negative adjustments (−10% or −20%), mostly being low-risk and low-income earners in the Basic scheme whose previous premium was not aligned with their risk level; 29% are unchanged, having already been nearly optimally priced under the baseline; and 40% of customers receive positive adjustments (+10% to +30%), corresponding to those who have multiple comorbidities or significant claim variance.

The reduction in the Gini coefficient from 0.412 to 0.331 shows the enhanced fairness in premium allocation in terms of different income groups, which aligns with the objectives of universal health coverage within Sustainable Development Goal 3. The graph below provides evidence of the improved equity by showing a comparison between the two Lorenz curves.

Table 5 Lorenz Curves Comparing Cumulative Premium Burden Distribution Across Income Deciles Under the Baseline Community-Rating System and the XAI-POS RL-Optimised Regime. The Narrowing Gap from the Line of Perfect Equality Reflects the Reduction in Gini Coefficient from 0.412 to 0.331.

Income Decile	Baseline Cumulative Premium Share (%)	XAI-POS Optimised Cumulative Share (%)
10% (lowest)	3.2	5.8
20%	8.1	13.2
30%	14.6	21.5
40%	22.8	31.0
50%	33.5	41.8
60%	46.2	53.9
70%	60.1	66.4
80%	74.8	79.2
90%	88.3	90.1
100% (highest)	100.0	100.0

Table 6 Summarises RL Agent Training Convergence Over 500 Episodes.

Episode	Cumulative Reward	Status
1–50	−120 to −80	Exploration phase ($\epsilon = 1.0$)
51–150	−80 to −20	Learning phase, reward improving
151–250	−20 to +40	Accelerated improvement
251–350	+40 to +75	Near-convergence
351–387	+75 to +82	Convergence criteria met
388–500	+82 (stable)	Policy stable: <3% state-action change

Table 6. RL agent cumulative reward per episode during training (500 episodes). Convergence was confirmed at episode 387 when all three criteria were satisfied: reward stabilisation within 5% tolerance, Q-value variance below 0.01, and policy stability above 97% over 50 consecutive episodes.

V. DISCUSSION

➤ Contribution to Actuarial and AI Literature

The XAI-POS offers four unique contributions within the domain of Actuarial Science and Deep Learning applications. Firstly, the implementation of the hybrid LSTM-CNN model proves superior compared to GLMs for individual cost predictions in a developing country setting, improving upon R^2 scores by 37% from Ridge Regression and 9% from the highest performing standalone algorithm. Secondly, RL-based premium optimisation achieves sustainable, viable and just distribution in an actuarial sense, something not possible when using traditional GLM models. Thirdly, XAI-POS provides a novel triple-XAI explainability framework (SHAP, LIME, permutation importance) tailor-made for the insurance industry and overcoming the transparency challenges that have prevented the adoption of AI in the field before. Lastly, according to the author’s

research, XAI-POS is the first ever publicly disclosed study utilising AI-based premium optimisation in Sub-Saharan Africa’s health insurance market (Carrin and James, 2021).

➤ Interpretation of Key Findings

The 18.9% decrease in average OOP deficit experienced with XAI-POS on the holdout cohort (2023-2025) is because of the following two factors: (1) premium segmentation based on higher risks results in better adequacy of the insurance cover and a closer match between the maximum benefit payout and cost; and (2) low-risk clients pay lower premiums, ensuring affordability and preventing selection bias, and thereby breaking the spiral of market degradation caused by underpricing of catastrophic risks.

The increase in the actuarial loss ratio from 1.04 to 0.79 has a major impact on financial sustainability. The value above 1.0 means that the market operates under losses and cannot expand further due to coverage. RL-based approach enables dynamic modification of premiums and ensures that the market remains in the range of financial sustainability of 0.60-0.85 – which is impossible with yearly premium changes.

A drop in the Gini ratio from 0.412 to 0.331 has great significance for Zimbabwe considering the established connection between income inequality and disparities in health. The RL reward function punishes premiums above 10 percent of monthly income, thus making the pooling of premiums easier for subsidization of high-risk individuals.

➤ *Limitations*

These limitations need to be noted. Firstly, while the IPEC data of 1,500 records is statistically strong enough to develop and test the model, for implementation in reality, we will need the entire population of the IPEC policyholders. Secondly, the weights of the RL reward function, which are currently arbitrarily selected, need to be determined by participatory methods with involvement of IPEC, medical aid society members, and the general public before seeking regulatory approval. Thirdly, the model does not factor in behavioural reactions of the policyholders to changing premiums, which may impact the dynamics of adverse selection, and therefore, the inclusion of game-theoretic extensions to the RL model is suggested for future research. Finally, although the outcomes of this analysis were validated for 2023-2025, future application needs to consider economic circumstances which may be quite different from those seen during training and validation.

VI. POLICY RECOMMENDATIONS

➤ *Insurance and Pensions Commission of Zimbabwe (IPEC)*

- Creation of risk-adjusted premium levels using hybrid DL models, thereby superseding the existing community rating model.
- A legislative stipulation of loss ratio within the range of 0.60 and 0.85 based on AI-empowered actuarial models that account for macro-economic data on a quarterly basis.
- Formation of regional pools to ensure that the premium Gini index falls under 0.30 within the next three regulatory cycles.
- Introduction of OOP catastrophic expenses cover for all Basic policyholders based on the recommended 10% CHE level at WHO.
- Introduction of regulatory sandboxing of AI-empowered insurance premium pricing algorithms. The first of such programmes would be XAI-POS with LIME explanations offered to each policyholder separately.

➤ *Medical Aid Societies and Health Insurance Companies*

- Implement dynamic premium setting utilizing reinforcement learning algorithms, which factor in inflation, exchange rates, and claims trend analysis.
- Develop disease management programs for HIV/AIDS, diabetes, high blood pressure, and heart disease, validated by SHAP analysis as the primary causes of OOP expenses.

- Use LSTM time series modeling in place of retrospective claims analysis to perform actuarial reserving and future expense forecasting.
- Revise Basic plan benefits through convolutional neural network interaction analysis, enabling coverage against multiple risks based on individual risk factors.

➤ *Ministry of Health and Child Care*

- Creation of a standard health insurance package using AI's predicted utilisation patterns based on geographical and demographic segments.
- Concentration on basic and preventative health insurance packages that have been shown by XAI to have higher utilisation and lower costs per episode.
- Rural health facility development in those provinces where the provincial heat map indicates heavy dependence on out-of-pocket spending because of lack of specialists.
- Using the NSSA contributions system to cover the costs of health insurance premium payments for formal sector workers.

➤ *Ministry of Finance*

- Craft premium tax incentives for people who have more than the 5% threshold set by WHO with regard to premium income, using the XAI-POS affordability index. This is estimated to cover around 38% of currently insured people.
- Maintain healthcare financing stability reserves through excise taxes on tobacco and alcohol in order to prevent the devaluation of the insurance medical coverage due to inflationary effects.
- Link minimum regulated premiums to the healthcare consumer price index, which is tied to economic growth, and thereby eliminating the gap between premiums and claims which accounts for 37% of OOP expenses.
- Establish a premium matching scheme to attract private sector funding into rural health facilities.

VII. CONCLUSION

XAI-POS represents a novel explainable deep learning premium optimization system with LSTM, CNN, and RL models for medical insurance in Zimbabwe. This is trained using a temporally split data of 1,500 IPEC policies (with training period of 2020–2022 and test set of 2023–2025) and convergence after episode 387 of 500 training episodes, giving an out-of-sample individual cost prediction R^2 of 0.88 ± 0.02 . Furthermore, it lowers the mean OOP deficit by 18.9%, stops 11.5 percentage points of overall medical expenditure from being tragically, improves the premium affordability index from 58 to 74, lowers the actuarial loss ratio from an unsustainably high 1.04 to 0.79, and lowers the premium Gini coefficient from 0.412 to 0.331. As explained above, with the explainability provided through SHAP, LIME, and permutation importance, all the premium recommendations become completely clear and can be audited by IPEC, insurers, and individual policyholders,

eliminating the black-box issue which has prevented AI deployment in financial services before.

Here are three suggestions for future research directions: (i) incorporation of game-theoretic modelling of policyholder behavior in the RL model; (ii) validation of the model against the whole population of IPEC policyholders; and (iii) implementation of real-time deployment tool for actuaries and IPEC auditors.

ACKNOWLEDGEMENTS

The authors would like to express their gratitude to the Insurance and Pensions Commission of Zimbabwe (IPEC) for providing the data in agreement number IPEC-DAA-2024-009. The authors further thank the Research Ethics Committee of the Harare Institute of Technology for the ethical approval accorded in agreement HIT-REC-2024-0147. This work is the personal viewpoint of the authors and may not be taken to reflect the views of IPEC or HIT.

REFERENCES

- [1]. Abiodun, O.I., Jantan, A., Omolara, A.E., Dada, K.V., Mohamed, N.A. and Arshad, H. (2021). State-of-the-art in artificial neural network applications: A survey. *Heliyon*, 4(11), e00938.
- [2]. Braunsteins, P. and Mandjes, M. (2023). Advances in actuarial science: A critical review of traditional model limitations in modern healthcare contexts. *Journal of Risk and Insurance*, 90(2), pp. 341–378.
- [3]. Carrin, R. and James, C. (2021). Social health insurance: Key factors affecting the transition towards universal coverage. *International Social Security Review*, 74(1), pp. 23–48.
- [4]. Chikanda, A. (2021). Medical leave: The exodus of health professionals from Zimbabwe. *Journal of Southern African Studies*, 32(3), pp. 507–526.
- [5]. Dlamini, S. and Sibanda, M. (2023). Machine learning for medical aid claims prediction in eSwatini. *African Journal of Health Economics*, 12(1), pp. 45–62.
- [6]. Frees, E.W. (2022). *Regression Modelling with Actuarial and Financial Applications*. 2nd ed. Cambridge: Cambridge University Press.
- [7]. Goodfellow, I., Bengio, Y. and Courville, A. (2016). *Deep Learning*. Cambridge, MA: MIT Press.
- [8]. Goodman, B. and Flaxman, S. (2017). European Union regulations on algorithmic decision-making and a 'right to explanation'. *AI Magazine*, 38(3), pp. 50–57.
- [9]. Henckaerts, R., Côté, M.P., Antonio, K. and Verbelen, R. (2022). Boosting insurance pricing with machine learning. *ASTIN Bulletin*, 52(1), pp. 1–30.
- [10]. Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), pp. 1735–1780.
- [11]. Insurance and Pensions Commission of Zimbabwe (IPEC) (2022). *Annual Insurance and Pensions Industry Report 2022*. Harare: IPEC.
- [12]. Insurance and Pensions Commission of Zimbabwe (IPEC) (2023). *Annual Insurance and Pensions Industry Report 2023*. Harare: IPEC.
- [13]. Insurance and Pensions Commission of Zimbabwe (IPEC) (2024). *Data Access Agreement IPEC-DAA-2024-009: Administrative Returns Dataset for Academic Research*. Harare: IPEC.
- [14]. Insurance and Pensions Commission of Zimbabwe (IPEC) (2025). *Field Consultation on Health Insurance Shortfall Challenges in Zimbabwe*. [Unpublished field report]. Harare: IPEC.
- [15]. James, G., Witten, D., Hastie, T. and Tibshirani, R. (2021). *An Introduction to Statistical Learning with Applications in R*. 2nd ed. New York: Springer.
- [16]. Kaas, R., Goovaerts, M., Dhaene, J. and Denuit, M. (2021). *Modern Actuarial Risk Theory: Using R*. 2nd ed. Hoboken, NJ: Wiley.
- [17]. Kuo, K. and Richman, R. (2021). Convolutional neural networks for insurance pricing. *ASTIN Bulletin*, 51(2), pp. 453–490.
- [18]. Lundberg, S.M., Erion, G., Chen, H., DeGrave, A., Prutkin, J.M., Nair, B. and Lee, S.I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1), pp. 56–67.
- [19]. Naik, N., Hameed, B.M.Z., Shetty, D.K., Swain, D., Shah, M., Paul, R. and Somani, B.K. (2022). Legal and ethical consideration in artificial intelligence in healthcare: Who takes responsibility? *Frontiers in Surgery*, 9, 862322.
- [20]. Rajpurkar, P., Chen, E. and Lungren, M.P. (2022). AI in health and medicine. *Nature Medicine*, 28(1), pp. 31–38.
- [21]. Ratner, B. (2023). *Statistical Modelling and Analysis for Database Marketing: Effective Techniques for Mining Big Data*. Boca Raton: CRC Press.
- [22]. Reserve Bank of Zimbabwe (RBZ) (2024). *Monetary Policy Committee Statement and Financial Stability Report*. Harare: RBZ.
- [23]. Richman, R. (2021). AI in actuarial science — A review of recent advances. *Annals of Actuarial Science*, 15(2), pp. 207–229.
- [24]. Wagstaff, A. and Neelsen, S. (2020). A comprehensive assessment of universal health coverage in 111 countries. *npj Global Health*, 3(1), p. 22.
- [25]. Wiley, J. (2021). Multi-objective optimisation in insurance pricing: Pareto frontier approaches. *Insurance: Mathematics and Economics*, 98, pp. 114–129.
- [26]. World Health Organisation (WHO) (2022). *Global Health Expenditure Database*. Geneva: WHO. Available at: <https://apps.who.int/nha/database>.
- [27]. Wüthrich, M.V. and Merz, M. (2023). *Statistical Foundations of Actuarial Learning and Its Applications*. Cham: Springer.
- [28]. Zimbabwe National Statistics Agency (ZimStat) (2020). *Zimbabwe Poverty Update 2020: Household Income and Expenditure Survey*. Harare: ZimStat.