

Multi-Modal Data Fusion for Life and Annuity Risk Modelling

Sita Rama Praveen Madugula
Independent Researcher/Technology Lead

Abstract—The life and annuity insurance sector is undergoing a profound digital transformation as data sources proliferate across structured policy records, unstructured narratives, wearable devices, and macroeconomic indicators. Traditional actuarial models, while robust, are increasingly limited in capturing the full spectrum of risk embedded in these heterogeneous datasets. Multi-modal data fusion offers a systematic approach to integrate such diverse inputs into unified machine learning models, enabling richer insights into underwriting, fraud detection, and long-term risk assessment. This survey reviews the current state of multi-modal fusion architectures, including early, intermediate, late, and hybrid strategies, and evaluates their applicability to life and annuity risk modelling. It further examines the role of data harmonisation, interpretability, and regulatory compliance as both enablers and constraints of adoption. Case applications are illustrated through underwriting enhancements, fraud detection improvements, and capital requirement modelling. The study concludes by addressing outstanding challenges in explainability, ethical governance, and data quality, while highlighting opportunities presented by emerging paradigms such as federated learning, generative modelling, and hybrid actuarial–AI systems.

Index Terms—Multi-modal data fusion; Life insurance; Annuity modelling; Machine learning; Underwriting; Fraud detection; Risk assessment; Explainable AI; Data harmonisation.

I INTRODUCTION

1.1 MOTIVATION AND CONTEXT

Insurance has long been characterised by its reliance on structured data such as age, gender, income, and medical records to assess the financial risks of mortality and longevity. However, the modern digital ecosystem has introduced an unprecedented volume of new modalities. Wearable devices monitor physical activity and physiological markers in real time, text

documents contain nuanced medical and claims information, while macroeconomic indicators capture systemic exposures. Harnessing these diverse sources is no longer a matter of incremental innovation but a necessity for insurers seeking competitiveness, solvency, and regulatory compliance in a volatile market.

Table I. Heterogeneous Data Sources in Life and Annuity Risk Modelling

Structured Policy	Age, income, premiums, duration	Traditional underwriting and capital modelling	Limited contextual variability
Unstructured Text	Claims notes, medical narratives	Detects anomalies, richer context for underwriting	Complexity of NLP, semantic bias
Wearables	Heart rate, steps, sleep, oxygen levels	Real-time health monitoring, dynamic longevity assessment	Data quality, privacy concerns
Macroeconomic	GDP, inflation, unemployment	Captures systemic risks affecting solvency	Volatility, correlation shifts

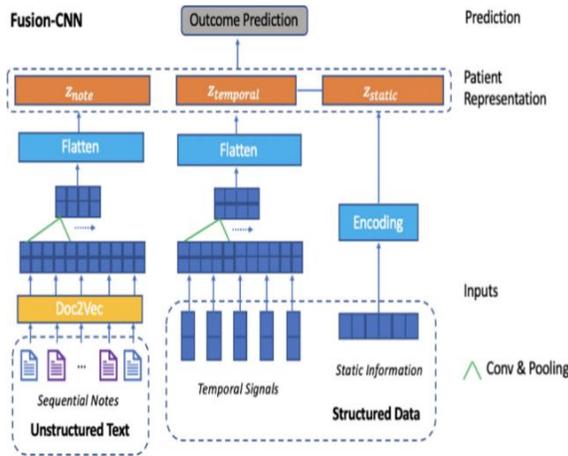


Figure I: High-level schematic of data modalities

A Fusion-CNN architecture integrating unstructured text, temporal signals, and structured data. Each modality is encoded separately before being fused into a joint representation for outcome prediction.

1.2 FROM ACTUARIAL SCIENCE TO MACHINE LEARNING INTEGRATION

The conventional actuarial science uses deterministic and statistical methods and generates survival graphs and annuity payout projections that follow fairly close relationships. By contrast, the success of machine learning relies on an abundance of data and non-linearity, so it can learn complex dependencies in time and context. It is, in other words, a paradigm shift: machine learning and actuarial reasoning are becoming more personalised, more effective at detecting fraud, and more flexible in the appointment of capital.

II THE PURPOSE OF THE SURVEY

This paper provides a systematic review of multi-modal data fusion for life and annuity risk modelling. In particular, it discusses (i) the historical development of data modalities as they apply to insurance, (ii) the key fusion architectures and their technical underpinnings, (iii) use-cases of the underwriting, (iv) detection of fraud, and (iv) risk capital modelling, and (v) challenges and opportunities that will govern the future of the area. Through the transference of knowledge base in computer science, actuarial practice, and financial regulation, this study would serve two purposes: a guide on research as well as a stepping stone toward industrial implementation.

2.1 BACKGROUND AND PRELIMINARY

Products Arguably, the most relevant products of long-term financial planning are life insurance and annuities, whose essential role is to pass mortality and longevity risks to and from institutions and individuals. These contracts result in insurers having decades-long liabilities, and such an accurate risk forecast is necessary. The biometric models are more reliant on dynamic policyholder information and population-based mortality schedules that are not up to date with behaviour and environmental differences.

2.2 FORMATION OF MULTI-MODAL DATA FUSION

Multi-modal data fusion: The term multi-modal data fusion is used to denote the incorporation of different data modalities into combined data analytics. Initially debuting in the area of multimedia applications, its spread has reached areas like healthcare, finance and now even insurance. For life and annuity products, multi-modal fusion entails combining structured data (e.g., premiums, policy duration), unstructured narratives (e.g., medical notes, claims text), biometric time series (e.g., wearable data), and contextual indicators (e.g., inflation rates, employment indices). This approach offers a multi-layered view of risk.

2.3 RELATED PRECEDENTS IN ADJACENT FIELDS

Perhaps the most noteworthy precedent is provided by healthcare, where the combination of radiological imagery, genomics, and historical patient records has led to the earlier, more accurate detection of the presence of the disease. In financial services, a combination of metadata on transactions with textual and behavioural features has been used to help detect fraud. These analogies demonstrate the transformational (newspeak) powers of fusion techniques and show why the insurance setting is relevant to them.

III SURVEY METHODOLOGY

3.1 LITERATURE SELECTION

A systematic review protocol was followed in the methodology adopted. Google Scholar, IEEE Xplore, ScienceDirect, and SpringerLink have been used to acquire peer-reviewed sources. Other search terms included multi-modal fusion, life insurance, annuity risk modelling and fraud detection, and wearable health data. Both theoretical and empirical studies that

have been published between 2010 and 2024 were regarded.

3.2 INCLUSION AND EXCLUSION CRITERIA

Among these included were studies which addressed the following:

Fusion plans that can be applied to the structured and unstructured data,

Risk modelling with insurance data sets including non-homogeneous information

Explainability of machine learning in the regulated sector.

Excluded were:

Unimodal solutions that do not perform across modalities,

Non-financial applications that would not be transferable into insurance,

Theoretical explanations that are purely based on theory without computational or empirical complement. It is an application of a theory to the calculation, and the results satisfy certain requirements, such as the convergence in the limit.

3.3 GARDEN FESTERS SCOPE AND LIMITATIONS

The scope of the present survey is not nervous. This has been replaced with a more interdisciplinary scope where contributions of other related fields are deliberately included, and these include contributions of the healthcare field, contributions of financial fraud detection and risk analytics in banking. One can attribute the justification of this inclusivity to the fact that there are methodological similarities that can be traced in these domains. An example is healthcare, where multi-modal data fusion frameworks involving imaging, electronic health records, and genomic data have taken centre stage in their development. The same methods can be applied directly to the insurance industry, where data scientists and actuaries encounter a similar-at-root challenge of merging disparate modalities, including discrete structured policyholder data, unstructured medical histories, biometric data on wearable devices, and macroeconomic forces. Likewise, claim validation in insurance demands a similar set of cross-modal methods that combine structured transactional records with textual metadata or behavioural correspondence, as is required to detect financial fraud. By extending the scope of the survey to include these cross-domain results, this study gains a richer methodology and the assurance that practices that are state-of-the-art in more developed areas are

not ignored in the relative youth of the application area of insurance.

Meanwhile, the need to accommodate neighbouring-discipline literature points to the key limitation of the survey: the relative lack of empirical studies of life and annuity risk modelling. Large-scale, peer-reviewed comparisons of multi-modal systems have been made in health informatics and banking, but remain conceptual in the insurance industry, where there are limited longitudinal studies and mostly exploratory pilot studies or industry white papers. Accordingly, this paper must extend to some degree of speculation, since healthcare and financial methods may bear some relevance to the life and annuity industry.

Overall, the strategic expansion of the scope of this survey to the medical and allied professions contributes to the discussion and reinforces the methodological footing of the review. At the same time, however, it indicates the absence of research in insurance-related areas per se. A priority in future research should be to fill this gap, especially through interdisciplinary research between academia and industry and regulators. The promise of the use of multi-modality in refining life and annuity risk modelling cannot become a reality without such domain-specific validation.

IV ARCHITECTURES - MULTI-MODALITY FUSION

4.1 EARLY FUSION

Early fusion carries out a direct concatenation of raw features or features which have been pre-processed across different modalities into one long vector before the model training takes place. The method has a simple conceptual appeal, and the simplicity can make this method attractive because it allows the model to retrieve information on all sources at the same time without having to use multiple specialised encoders. Specifically, policyholder demographic data, textual embeddings of medical notes, and continuous signals of wearable devices can be combined in one unified feature representation, which is in turn fed into a predictive model, e.g., a multilayer perceptron.

Nevertheless, there are considerable challenges with this seemingly simple method. Early fusion is exposed

to the curse of dimensionality in that concatenating high-dimensional representations could easily lead to vectors whose sizes cannot be used easily. In addition, disparities in scale or distribution across modalities, i.e. continuous versus categorical insurance variables versus continuous biometric data, may bias the model and lead to degraded performance. Most of these problems require heavy preprocessing or feature scaling and dimensionality reduction, i.e. PCA or autoencoders. That said, early fusion is useful in applications where the data size is relatively small and where the interactions amongst modalities are to be interpreted, such as in insurance.

4.2 MEDIOD FUSION

Intermediate fusion would keep modality-specific encoders, such as transformers in the case of text, convolutional neural networks in the case of time-series data and MLPs in tabular representations, to be combined only at the latent representation level. This maintains domain-specific patterns and allows cross-modality interactions, hence extremely effective in the context of underwriting tasks where biometric signal, structured policy data, and economic indicators have to be combined manually. The primary difficulty is an upsurge in model complexity and a necessity to use bigger, high-quality datasets.

4.3 LATE FUSION

Late fusion combines results of models trained independently (e.g. actuarial survival curves can be used together with fraud-identification classifiers to produce composite-risk scores). The modularity of the system enables insurers to be transparent and regulatory-focused because each of the components can be validated individually. Since only representative fusion happens at the decision level, it does not capture intricate feature-level synergies across modalities, and thus, can become worthless in prediction than previous methods of integration.

4.4 HYBRID FUSION

Hybrid fusion combines early type strategies, intermediate type and late type strategies, enabling the possibility to interact at various abstraction levels. In modelling life and annuity, this allows the structured, unstructured and macro-economic data to be incorporated within the same analysis to improve solvency projections and fraud detection, respectively.

In spite of the flexibility and power of hybrid approaches, such methods are computationally demanding and may become black-box systems. However, used with explainable AI, they are one of the most promising avenues of advanced insurance risk modelling.

Table II. Comparison of Fusion Strategies in Insurance Applications

Early Fusion	Feature-level	Captures cross-modal interactions	High dimensionality	Policy + text embeddings
Intermediate Fusion	Latent representations	Balances accuracy and interpretability	Requires complex design	Wearable + macroeconomic fusion
Late Fusion	Decision outputs	Modular, regulator-friendly	Ignores feature-level detail	Actuarial + ML blending
Hybrid Fusion	Multi-level	Robust, flexible	Computationally demanding	Enterprise pipelines

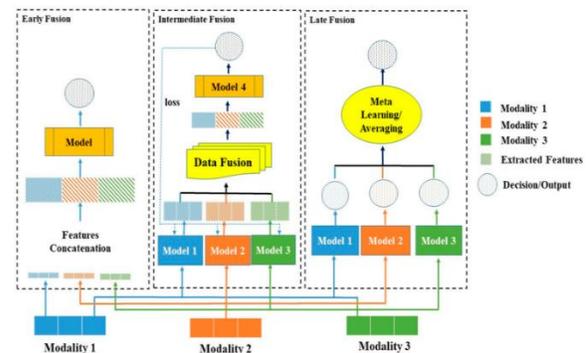


Figure II: Diagram illustrating early, intermediate, late, and hybrid fusion workflows applied to life and annuity risk modelling

This diagram shows the different fusion workflows—early, intermediate, late, and hybrid—used in life and annuity risk modelling. It highlights how data and model outputs can be combined at various stages of the process to improve accuracy and decision-making.

V APPLICATIONS IN LIFE AND ANNUITY RISK MODELLING

5.1 UNDERWRITING

Underwriting is the process of measuring risk (properly sizing risk), and has historically relied on age, gender, occupation, and medical history. This process is enriched by the addition of wearable health signals, such as a real-time index of heart rate, sleep quality values, or activity rates. Combined with structured application data and text-based clinical notes, insurers are able to produce highly detailed risk analysis. In this case, intermediate fusion models could be particularly well suited to do this, being the combination of physiological time series data with macroeconomic variables, a combination of personal and systemic aspects. This enables insurance companies to encompass not only current health risks, but also future impacts that economic fluctuations may have on future liabilities.

5.2 FRAUD DETECTION

Cheating normally occurs in the form of textual anomalies in the storyline, or inconsistency in action with regard to biometric data or history. The multi-modal fusion allows carriers to correlate the organised claims with auxiliary biometric sensors and written reports. As an example, the data provided in a medical claim of long-term immobility can be compared with a wearable step tracker to identify inconsistencies. A holistic system would greatly minimise false positives at a minimal cost to detection frameworks and provide a greater resilience against emerging fraud strategies.

5.3 RISK ASSESSMENT AND CAPITAL MODELLING

Life and annuity insurers should be in a position to carry out capital modelling accurately long-term estimation of liabilities. The Hybrid fusion is a solution with additional strength as it combines actuarial survival models and macroeconomic forecasts with policyholder behaviour data. This method enhances solvency capital requirement (SCR), therefore, helping the insurers resist systemic shocks. The combination of capturing micro-level behaviour and macroeconomic uncertainty allows hybrid models to provide a higher level of financial stability as well as more adaptive risk-based capital strategies.

Table III. Applications of Multi-Modal Fusion in Insurance

Domain	Modalities Integrated	Benefits
Underwriting	Policy data, wearables, text, macroeconomic	Personalised risk profiling
Fraud Detection	Claims data, text narratives, biometrics	Reduced false positives
Capital Modelling	Actuarial survival models, macroeconomic data, and behaviour	Enhanced solvency forecasting



Figure III: Workflow diagram of multi-modal integration applied to underwriting, fraud detection, and capital modelling.

This diagram above illustrates a workflow for multi-modal integration applied to underwriting, fraud detection, and capital modelling, showing how diverse data sources are combined to enhance accuracy, efficiency, and risk assessment.

VI CHALLENGES AND OPPORTUNITIES

6.1 INTERPRETABILITY

Complex systems can end up building a black-box system. Insurers need to use explainable AI techniques relevant to multi-modality fusion, which can include modality-specific Shapley values or hierarchical attention visualisations.

6.2 REGULATION AND ETHICS

This chapter is concerned with learning and knowing. It has increasingly become a mechanism that tries to practically or empirically accomplish much of what is currently Regulation and Ethics, these few topics in life. Regulation

Regulatory bodies require fairness, transparency and accountability. Fusion solutions that include sensitive biometrics will have to be designed to a high standard and with due consideration given to GDPR and Solvency II.

6.3 DATA HARMONISATION

Disparities between modalities, such as differences in sampling frequencies as well as missing values, act as a hindrance to the training of coherent models. State-of-the-art imputation and domain adaptation methods are important remedies.

6.4 PROSPECTIVE DIRECTIONS

New paradigms are promising. Federated learning can provide model training with multiple parties conducting it without compromising data privacy. Generative models have the power to generate synthetic data sets which are statistically strong and are used in cases of an imbalanced class. AI systems in the context of hybrid actuarial frameworks hold the most potential to become the affecting halfway point between understandable and powerful predictions.

6.5 CONCLUSION

This survey has shown that multi-modal data fusion is an important breakthrough with regard to the trend of life and annuity risk modelling. By harmonising structured policy in the records, unstructured text, biometrics wearables, and macroeconomic data, insurers will be able to create more precise, robust and personalised models of risk. The paper identified four main strategies of fusion: early, intermediate, late, and hybrid-and positioned them in the context of their application underwriting, fraud detection, and solvency modelling. The resulting analysis also emphasised the problems of interpretability, regulatory compliance, and data harmonisation as long-standing issues, but also highlighted potentially auspicious trends in federated learning, generative data augmentation, and hybrid actuarial-AI systems.

To sum up, multi-modal fusion in life and annuity contexts is a structural form of innovation and cannot be seen as tangential transformations but rather as those that redefine the field. To scientists, it implies further developing techniques that can come up with a compromise between the accuracy of prediction and

transparency. To practitioners, it involves the establishment of operational frameworks that would ethically and responsibly exploit heterogeneous data without breaching ethics or regulation.

REFERENCES

- [1] Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Benetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2019). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.07.007>
- [2] Awosika, T., Shukla, R. M., & Pranggono, B. (2023). Transparency and privacy: The role of explainable AI and federated learning in financial fraud detection. *arXiv preprint arXiv:2312.13334*. <https://doi.org/10.48550/arxiv.2312.13334>
- [3] Baldi, P., & Sadowski, P. (2014). The dropout learning algorithm. *Artificial Intelligence*, 210, 78–122. <https://doi.org/10.1016/j.artint.2013.12.001>
- [4] Baltrušaitis, T., Ahuja, C., & Morency, L.-P. (2017). Multimodal machine learning: A survey and taxonomy. *arXiv preprint arXiv:1705.09406*. <https://arxiv.org/abs/1705.09406>
- [5] Chen, L., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. <https://doi.org/10.1145/2939672.2939785>
- [6] Feng, R., & Volkmer, H. W. (2012). Analytical calculation of risk measures for variable annuity guarantee benefits. *Insurance: Mathematics and Economics*, 50(1), 1–15. <https://doi.org/10.1016/j.insmatheco.2011.10.001>
- [7] Feng, R., & Volkmer, H. W. (2014). Spectral methods for the calculation of risk measures for variable annuity guaranteed benefits. *ASTIN Bulletin: The Journal of the IAA*, 44(3), 653–681. <https://doi.org/10.1016/j.astin.2014.04.001>
- [8] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2672–2680. <https://arxiv.org/abs/1406.2661>

- [9] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2672–2680. <https://arxiv.org/abs/1406.2661>
- [10] Govindaswamy, R., & Subbian, R. (2023). Advanced data-driven frameworks for intelligent underwriting risk assessment. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 9(2), 880–893. <https://doi.org/10.32628/CSEIT2342437>
- [11] Groll, A., Khanna, A., & Zeldin, L. (2024). A machine learning-based anomaly detection framework in life insurance contracts. *arXiv preprint arXiv:2411.17495*. <https://doi.org/10.48550/arxiv.2411.17495>
- [12] He, J., Bai, Y., & Wang, W. (2019). Multimodal deep learning for health risk prediction. *IEEE Access*, 7, 87209–87219. <https://doi.org/10.1109/ACCESS.2019.2924724>
- [13] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [14] He, X., Zhang, H., Kan, M.-Y., & Chua, T.-S. (2016). Fast matrix factorization for online recommendation with implicit feedback. *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 549–558. <https://doi.org/10.1145/2911451.2911500>
- [15] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *CVPR Proceedings*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [16] Hejazi, S. A., & Jackson, K. R. (2016). A neural network approach to efficient valuation of large portfolios of variable annuities. *arXiv preprint arXiv:1606.07831*. <https://arxiv.org/abs/1606.07831>
- [17] Huang, F., Shen, J., Yang, Y., & Zhao, R. (2024). Fairness-aware principal component analysis for mortality forecasting and annuity pricing. *arXiv preprint arXiv:2412.04663*. <https://doi.org/10.48550/arxiv.2412.04663>
- [18] Kim, J., & Rush, A. M. (2016). Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947*. <https://arxiv.org/abs/1606.07947>
- [19] Liu, B. (2021). Robust dynamic multi-modal data fusion: A model uncertainty perspective. *arXiv preprint arXiv:2105.06018*. <https://doi.org/10.48550/arxiv.2105.06018>
- [20] Shankar, S., Thompson, L., & Fiterau, M. (2022). Progressive fusion for multimodal integration. *arXiv preprint arXiv:2209.00302*. <https://doi.org/10.48550/arxiv.2209.00302>
- [21] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008. <https://arxiv.org/abs/1706.03762>
- [22] Yang, J., Chen, K., Ding, K., Na, C., & Wang, M. (2022). Auto insurance fraud detection with multimodal learning. *Data Intelligence*, 5(2), 388–412. https://doi.org/10.1162/dint_a_00191
- [23] Zhang, D., Nayak, R., & Bashar, M. A. (2021). Exploring fusion strategies in deep learning models for multi-modal classification. In *Proceedings of the Australasian Conference on Data Mining (AusDM 2021)* (pp. 102–117). Springer. https://doi.org/10.1007/978-981-16-8531-6_8
- [24] Zhang, Z., Chen, M., & Wang, F. (2020). Multimodal fusion with deep neural networks for financial risk assessment. *Knowledge-Based Systems*, 195, 105644. <https://doi.org/10.1016/j.knosys.2020.105644>
- [25] Zhou, X., & Zhang, Y. (2022). Analysis of multimodal data fusion from an information theory perspective. *Information Sciences*, 585, 1–17. <https://doi.org/10.1016/j.ins.2022.04.026>