

Research_Paper

INSURANCE MANAGEMENT SYSTEM

An AI-Assisted Approach to Policy Management and Claims Fraud Flagging

Author: Anjali Sharma **Programme:** BCA — Final Year Project, AI Foundation Programme

Institution: SJ Global PU & Degree College, in collaboration with Aider Digital Pvt Ltd

Abstract

Insurance management involves activities such as policy administration, premium calculation, claims submission, claims assessment, customer servicing, and maintenance of policy records. The increasing digitisation of insurance has moved many of these activities from paper-based and manually coordinated processes toward digital platforms that can store, process, and retrieve information more efficiently. Digital technologies are also being applied across the insurance value chain, including underwriting, pricing, claims management, customer service, and fraud detection.

One particularly important problem is insurance claim fraud. Fraudulent claims create financial costs for insurers and policyholders, while identifying suspicious claims can itself be time- and resource-intensive. A 2023 study in the Journal of Risk and Insurance reported that insurance fraud detection is complex and costly and evaluated both supervised and unsupervised machine-learning approaches using real automobile insurance claims. Earlier research has similarly demonstrated the use of neural networks, Bayesian methods, statistical classification, and anomaly-detection techniques for insurance fraud detection.

This paper proposes a realistically scoped **Insurance Management System** for a final-year BCA project. The system provides digital management of insurance plans, policies, premium calculations, claims submission, claims tracking, and administrative review. Its AI component is deliberately limited to **fraud flagging and risk scoring**: a machine-learning model may assign a claim a suspiciousness score or flag based on patterns in the available data, but it does not automatically approve, reject, settle, or deny a claim. This design reflects research and regulatory guidance emphasizing transparency, accountability, record keeping, explainability, and human oversight in consequential insurance AI applications. The proposed system therefore treats AI as a **decision-support tool for human reviewers**, rather than as an autonomous decision-maker.

1. Introduction

Insurance is fundamentally a data-driven activity because insurers collect and analyse information to assess risk, determine premiums, manage policies, process claims, and identify possible fraud. Historically, many insurance activities depended on physical documents, manual data entry, human review, and communication between customers, agents, claims personnel, and administrative

departments. The availability of digital platforms has progressively changed this model by enabling electronic storage, digital distribution, online customer services, and automated processing of parts of the insurance value chain.

The transition toward digital insurance is not simply a technological change; it also changes how information is collected, accessed, analysed, and used. EIOPA has described digitalisation as increasingly important for the design, development, and distribution of insurance products and services through digital platforms and channels. EIOPA has also reported that data analytics has traditionally supported underwriting, pricing, claims settlement, and fraud detection, while the volume and variety of available data have expanded substantially through technologies such as connected devices and digital platforms.

Claims processing is an especially important area for digitalisation because claims require the insurer to collect information, verify policy coverage, assess the circumstances of the loss, and determine the appropriate next step. Digital platforms can provide structured claim submission, status tracking, document management, and faster access to policy information. EIOPA has identified faster processes, including claims-management processes, as one of the opportunities associated with financial innovation and digitalisation.

At the same time, insurance fraud remains an important research problem. Fraud detection can require substantial human effort because only a relatively small proportion of claims may ultimately be confirmed as fraudulent, while some suspicious claims may remain unidentified. A study by Debener, Heinke, and Kriebel used 7,750 automobile insurance claims from a German insurer and found that both supervised and unsupervised machine-learning methods could identify claims associated with fraud, while also demonstrating that the two approaches could identify different cases.

These developments provide a rationale for an insurance management system that combines ordinary digital management functions with a limited AI-assisted fraud-screening capability. The proposed system does not attempt to reproduce the complexity of a commercial insurance platform. Instead, it focuses on a small number of core functions that can realistically be designed, implemented, tested, and evaluated by one BCA student during one academic semester.

2. Background

2.1 From Manual Insurance Administration to Digital Platforms

Insurance has historically relied heavily on data and statistical analysis, but the cost and difficulty of collecting and processing information placed practical limits on the amount of information that could be used. EIOPA's 2021 report notes that insurance has long used mathematical and statistical methods for underwriting risks, pricing policies, quantifying losses, paying claims, and identifying or preventing fraud.

Digitalisation has changed the technological environment in which these activities occur. EIOPA describes the insurance sector as undergoing digital transformation involving digital platforms, ecosystems, artificial intelligence, open finance, blockchain, and Internet-of-Things technologies.

Digitalisation can affect product design, distribution, underwriting, customer service, and claims management rather than being restricted to a single part of the insurance business.

The movement toward digital insurance also creates practical opportunities for centralised policy management. Instead of maintaining separate paper records for policy information, payments, and claims, a digital system can maintain related records in a database and provide controlled access to different users. A customer can view policy information, an agent can assist with applications, and an administrator can review claims and payments through role-specific interfaces.

Digitalisation, however, does not remove the need for human judgement. Insurance remains a regulated activity, and the use of technology introduces additional requirements relating to data quality, security, fairness, accountability, and governance. Therefore, a student project should treat digitalisation as a means of improving information management and workflow rather than as a justification for replacing professional insurance judgement.

2.2 Insurance Fraud as a Research Problem

Insurance fraud is a well-established subject of academic research, particularly in automobile and health insurance. Fraud may involve deliberately fabricated losses or claims, or manipulation and inflation of otherwise legitimate claims. Academic literature has examined statistical classification, neural networks, anomaly detection, and other machine-learning methods for identifying suspicious claims.

The economic significance of the problem is also documented in the research literature. Debener, Heinke, and Kriebel cite an Insurance Europe estimate that fraudulent claims in Europe amounted to approximately €13 billion in 2017. This figure refers specifically to the cited 2017 estimate and should not be interpreted as a current global estimate of insurance fraud.

Fraud detection is difficult because confirmed fraudulent claims may represent only a small portion of all claims. In Debener, Heinke, and Kriebel's dataset, only 34 claims were ultimately identified as fraudulent among 7,750 automobile claims, while 162 claims had been marked as highly suspicious. The authors emphasised that actual fraudulent cases may be higher because some fraud may remain unidentified.

This characteristic creates a class-imbalance and labelling problem for machine-learning systems. Supervised learning requires labelled examples, but insurance organisations may have relatively few confirmed fraud cases and may contain fraudulent cases that have never been detected. Consequently, a model trained only on historical confirmed fraud can reproduce existing detection patterns and potentially miss previously unknown forms of fraud.

For these reasons, an appropriate student system should not claim that an AI model can determine whether a claim is fraudulent with certainty. A more defensible objective is to use machine learning to **prioritise claims for human review**.

3. Literature Review

3.1 Early Machine-Learning Research in Insurance Fraud

Research into computational insurance fraud detection predates current AI systems. Viaene, Dedene, and Derrig investigated Bayesian-learning neural networks for automobile personal-injury-protection claim fraud using 1,399 closed claims from Massachusetts accidents in 1993. Their research examined the ability of neural-network classifiers to identify relevant inputs and improve the interpretability of fraud classification.

The study is significant because it illustrates two issues that remain important in contemporary insurance AI: predictive performance and interpretability. The authors specifically examined which inputs were informative to the model rather than treating the neural network solely as a prediction mechanism.

Caruana and Grech later compared artificial neural networks and Naïve Bayes classification for automobile insurance fraud detection. Their study highlighted the fact that fraudulent claims constitute a relatively small portion of the dataset and therefore create an imbalanced-classification problem. Their results also demonstrated that model performance can vary between training and test data, reinforcing the importance of evaluating models on data that are not simply the same observations used for model development.

These studies support the conclusion that machine learning can be useful for insurance fraud research, but they do not justify treating a model's output as proof of fraud.

3.2 Unsupervised Learning and Anomaly Detection

Gomes, Jin, and Yang examined insurance fraud detection using unsupervised deep learning. Their work proposed the use of autoencoders and variational autoencoders together with a variable-importance methodology to obtain insights into potentially fraudulent behaviour without relying entirely on labelled fraud observations.

Unsupervised approaches are relevant because confirmed fraud labels can be scarce. Debener, Heinke, and Kriebel subsequently compared supervised and unsupervised machine learning using proprietary automobile insurance claims. Their study used isolation forests as an unsupervised method and XGBoost as a supervised method.

The study found that isolation forests showed strong predictive performance for insurance claim fraud and that supervised XGBoost also performed strongly despite the limited number of labelled fraudulent claims. The authors further found that the methods identified different cases and concluded that supervised and unsupervised approaches should be regarded as complements rather than substitutes.

For a BCA project, this finding suggests that a simple anomaly-detection or supervised classification model can be used as a **screening mechanism**, provided that its output is presented as a risk indicator rather than a definitive finding.

3.3 Machine Learning with Feature Selection and Classification

Aslam, Hunjra, Ftiti, Louhichi, and Shams studied automobile insurance fraud detection using machine-learning models and feature selection. Their framework used the Boruta feature-selection algorithm and compared logistic regression, support-vector machines, and Naïve Bayes. The researchers reported that support-vector machines performed best according to accuracy in their experiments, while logistic regression achieved the highest F-measure.

This study demonstrates that model selection should not be based on a single metric. Accuracy, precision, recall, F-measure, and other evaluation measures can produce different conclusions about the usefulness of a model. This is particularly important for fraud detection because a dataset can be highly imbalanced, making raw accuracy potentially misleading.

A student implementation should therefore report multiple evaluation measures where suitable, especially precision, recall, F1-score, and a confusion matrix.

3.4 Explainability in Insurance Pricing and Risk Assessment

Machine learning has also been studied for insurance pricing and risk classification. Kuo and Lupton examined explainability of machine-learning models in insurance pricing and argued that the regulated nature of insurance pricing creates a need for transparency in models. Their work presents techniques such as variable importance, partial-dependence analysis, and other model-explanation approaches for helping stakeholders understand machine-learning models.

Xie similarly investigated explainability of artificial neural networks for automobile insurance rate regulation and emphasised the importance of identifying major risk factors when machine-learning methods are used in insurance pricing. These studies indicate that predictive performance alone is insufficient when a model is used in a regulated insurance context.

For the proposed student project, this literature supports keeping the premium calculation component **non-AI and transparent**. A simple documented premium formula can be implemented without pretending that a student-built model is suitable for determining real insurance prices.

4. Regulatory and Ethical Context

4.1 Why Human Oversight Matters

The use of AI in insurance raises concerns beyond ordinary software development because algorithmic outputs can affect consumers' access to insurance, prices, claims handling, and other important outcomes. EIOPA's governance principles identify fairness, transparency and explainability, human oversight, data governance and record keeping, and robustness and performance as important principles for trustworthy AI in insurance.

EIOPA specifically states that adequate human oversight is important for AI-based fraud detection because fraud must ultimately be established by the insurance firm rather than assumed from an algorithmic output. This principle directly supports the design adopted in the proposed system: an AI model can flag a claim for investigation, but it cannot declare the claimant fraudulent.

EIOPA's 2025 Opinion on AI governance and risk management similarly emphasises data governance, record keeping, fairness, cybersecurity, explainability, and human oversight in the use of AI within insurance. The Opinion adopts a risk-based and proportionate approach rather than assuming that every AI application has identical risks.

4.2 European Union AI Act

The European Union's Artificial Intelligence Act, Regulation (EU) 2024/1689, provides an important example of a risk-based regulatory framework for AI. Under the Act, AI systems intended for risk assessment and pricing in life and health insurance fall within a specified high-risk category under Annex III.

The Act also establishes requirements concerning human oversight, transparency, and record keeping for high-risk AI systems. Article 14 requires high-risk AI systems to be designed so that natural persons can effectively oversee their operation and, where appropriate, interpret, disregard, override, or reverse AI outputs. Article 13 requires sufficient transparency to enable deployers to interpret high-risk AI system outputs and use them appropriately. Article 12 requires high-risk AI systems to technically allow automatic recording of relevant events to support traceability and monitoring.

An important qualification is necessary: **the EU AI Act does not classify every insurance claims-management AI system as high-risk under the life-and-health insurance provision.** The European Commission's AI Act guidance specifically distinguishes claims-management use cases from the Annex III category concerning life and health insurance risk assessment and pricing. Nevertheless, this does not mean that insurance claims AI is unregulated or that autonomous claims decisions are automatically appropriate. Existing insurance-sector governance frameworks continue to emphasise responsible use, data governance, transparency, and human oversight.

4.3 United States Insurance Regulatory Framework

The National Association of Insurance Commissioners (NAIC) adopted its Model Bulletin on the Use of Artificial Intelligence Systems by Insurers in December 2023. The bulletin states that decisions or actions affecting consumers that are made or supported by AI and related advanced analytical technologies must comply with applicable insurance laws and regulations.

The NAIC framework addresses risks including inaccurate outputs, unfair bias, discrimination, data vulnerabilities, governance, and risk management. It also expects insurers to maintain governance processes and documentation concerning the use of AI. The NAIC's current AI guidance further states that AI is expected to support rather than completely replace insurance professionals and that human oversight remains important in insurance decision-making.

Together, these regulatory and supervisory frameworks demonstrate why the proposed system should not be presented as an autonomous insurer. AI should provide information that assists a trained human reviewer while maintaining records of the AI output and the subsequent human decision.

5. Proposed System

5.1 System Overview

The proposed **AI-Assisted Insurance Management System** is a web-based application designed to centralise basic insurance operations while providing a limited machine-learning fraud-screening capability.

The proposed system contains four major functional areas:

1. Insurance policy management.
2. Premium calculation.
3. Claims submission and tracking.
4. AI-assisted claim fraud flagging and risk scoring.

The system can provide different interfaces for customers, insurance agents, and administrators. The customer can view insurance plans, submit applications, view policies, record premium payments, submit claims, and track claim status. Agents can assist with customer and policy operations, while administrators can review applications and claims.

The system is intended as an academic prototype rather than a production insurance platform. Its purpose is to demonstrate how a conventional information-management system can incorporate a limited AI decision-support component while preserving human responsibility.

5.2 Policy Management

The policy-management module maintains digital records of customers, insurance plans, applications, policy numbers, coverage, premium amounts, policy periods, and policy status.

A customer can select an available insurance plan and submit an application. The application can initially be marked as **Pending** and reviewed by an authorised administrator or agent.

The system can maintain policy states such as:

Pending → Active / Rejected → Expired

The system should maintain a history of policy-related actions so that changes can be traced.

5.3 Premium Calculation

Premium calculation should remain deliberately simple and transparent in the student implementation.

A deterministic formula can consider factors such as:

- insurance type;
- coverage amount;
- policy duration; and
- predefined age or plan categories.

The system should show the user how the estimated premium was calculated rather than using an opaque machine-learning model.

This design is supported by the literature on explainability in insurance pricing, which identifies transparency as an important consideration when predictive models influence insurance rates.

The proposed project therefore does **not** claim that its premium formula represents actuarially valid real-world insurance pricing.

5.4 Claims Submission and Tracking

Customers can submit claims through a digital form containing information such as:

- policy number;
- claim type;
- claimed amount;
- incident date;
- description of the incident; and
- supporting information.

The claim receives a unique claim number and moves through a simple workflow:

Submitted → Under Review → Approved / Rejected → Settled

The workflow is controlled by authorised human users.

The system can display the current claim status to the customer, while administrators can view the claim information required for review.

5.5 AI-Assisted Fraud Flagging

The proposed AI component is a **fraud-risk flagging system**, not an autonomous fraud-decision system.

The model receives selected structured claim features, such as claim amount, policy information, claim timing, previous claim history, and other appropriate features available in the project dataset. The exact features should depend on the quality and availability of the dataset.

The model produces an output such as:

Low Risk — 0.18

Medium Risk — 0.54

High Risk — 0.87

or an equivalent fraud-risk score.

The score means that the claim contains patterns that the model associates with previously observed suspicious cases. It does **not** mean that the claimant has committed fraud.

The administrator can therefore see:


AI Fraud Risk: High — Human Review Required

The final claim decision remains with the authorised human reviewer.

This approach is consistent with research showing that machine-learning fraud scores can help identify claims for investigation and with EIOPA's emphasis on human oversight in insurance fraud detection.

5.6 Human-in-the-Loop Workflow

The proposed workflow is:



The AI model cannot directly change the claim to **Approved**, **Rejected**, or **Settled**.

This restriction is central to the proposed design because insurance AI research and regulatory guidance emphasise the need for accountability, explainability, and human oversight in consequential applications.

5.7 Explainability and Audit Trail

The proposed system should store:


- claim identifier;
- model version;
- risk score;
- risk category;
- important input factors;
- date and time
- human reviewer;

Such records would allow the student to demonstrate how the AI output was used without claiming that the model itself made the final decision.

This design also reflects the importance of record keeping and traceability identified in the EU AI Act and EIOPA's insurance AI governance principles.

6. System Architecture

The proposed application can be implemented using a conventional three-layer web architecture.



The frontend can provide role-specific interfaces, while Django manages authentication, business rules, database operations, and API endpoints.

The database stores policy, customer, payment, and claim information.

The machine-learning component operates as a separate analytical service or module and returns a risk score to the claims-management workflow.

This architecture is appropriate for a student project because it separates ordinary information management from the experimental AI component without requiring complex distributed infrastructure.

7. Discussion

7.1 Data Quality

The quality of an insurance AI system depends strongly on the quality and representativeness of the data used for development and evaluation. EIOPA's AI governance principles specifically state that data used in AI systems should be accurate, complete, appropriate, securely stored, and subject to appropriate governance.

Insurance fraud datasets also present a particular challenge because confirmed fraudulent cases may be rare and some fraudulent cases may never be identified. This can result in class imbalance and incomplete labels.

Consequently, a student-developed model trained on a small or synthetic dataset cannot be assumed to represent the behaviour of real insurance claims.

7.2 Model Performance

A high classification accuracy does not necessarily mean that a fraud-detection system is useful. Insurance fraud datasets can be imbalanced, meaning that a model could achieve high accuracy by predicting the majority class while failing to identify actual fraudulent claims.

The proposed project should therefore evaluate the model using measures such as precision, recall, F1-score, and a confusion matrix rather than reporting accuracy alone.

The model should also be tested on data that were not used during training. Research in insurance fraud detection demonstrates the importance of evaluating model behaviour on held-out data and, where possible, on previously unsuspected claims.

7.3 Explainability

Explainability is a significant issue because a claims reviewer must be able to understand why a case was flagged sufficiently to use the output responsibly. EIOPA recommends explainable AI approaches, particularly for high-impact insurance applications, while also requiring appropriate governance and accountability.

The proposed student system should therefore prefer a relatively interpretable model, such as logistic regression or a decision tree, or use suitable explanation techniques when a more complex model is selected.

The objective is not to claim that an explanation makes an AI system automatically trustworthy. Rather, explanation should help the human reviewer understand the basis and limitations of the model output.

7.4 Human Oversight

The proposed system deliberately separates **prediction** from **decision**.

The AI model predicts or flags.

The human reviewer decides.

This distinction is important because EIOPA's insurance AI governance principles identify human oversight as a core governance principle and specifically emphasise human oversight in fraud-detection applications. The EU AI Act likewise requires effective human oversight for high-risk AI systems and gives human operators the ability to interpret, disregard, override, or reverse outputs.

The system therefore should never display language such as:

"Fraud confirmed by AI."

Instead, it should display:

"High fraud-risk score — human investigation required."

This distinction prevents the project from incorrectly representing a statistical prediction as a factual or legal determination.

7.5 Scope of Project

A real insurance platform involves complex actuarial models, regulatory requirements, large datasets, information-security controls, integrations with external systems, claims investigation processes, and extensive testing. The proposed student system cannot reproduce these characteristics within a single semester.

The appropriate academic objective is therefore to demonstrate the **concept and architecture** of AI-assisted insurance management rather than claim production-level insurance capability.

A realistic one-semester implementation can provide:

- customer and user management;
- insurance-plan management;
- basic premium calculation;
- policy applications;
- policy records;
- premium-payment records;
- claim submission;
- claim tracking;
- administrator review;
- an experimental fraud-risk model;
- human review of AI flags;
- basic audit logging; and
- dashboard reporting.

The project should explicitly identify the AI model as a prototype and should not claim that its performance is suitable for real insurance deployment.

7.6 Ethical Scope

The system should avoid unnecessary sensitive personal information, use only data required for the project, and restrict access according to user roles.

The AI model should also avoid being presented as an objective measure of a person's honesty or character. A high score indicates model-predicted similarity to patterns associated with suspicious claims; it does not establish wrongdoing.

This distinction is important because EIOPA's governance framework emphasises fairness, non-discrimination, transparency, human autonomy, and appropriate data governance in insurance AI.

8. Limitations

The proposed system has several important limitations.

First, the quality of the AI component depends on the availability, accuracy, completeness, and representativeness of training data. A small academic dataset cannot capture all types of insurance fraud encountered in commercial practice.

Second, fraud labels can be incomplete because some suspicious claims may never be investigated or proven fraudulent. This creates a fundamental limitation for supervised learning.

Third, machine-learning performance may change when the model encounters claims that differ from its training data. Debener, Heinke, and Kriebel found that supervised and unsupervised approaches could identify different types of suspicious claims, illustrating that no single modelling approach should be assumed to detect every fraud pattern.

Fourth, explainability remains a challenge for complex machine-learning systems. Research in insurance pricing has specifically identified the difficulty of adopting less interpretable machine-learning approaches in regulated insurance contexts.

Fifth, the proposed premium calculation is an educational approximation and is not an actuarial pricing model. Real insurance pricing involves substantially more data, validation, regulatory requirements, actuarial expertise, and business constraints than a one-semester student application can reproduce.

Finally, the proposed AI component is not suitable for autonomous claim approval or rejection. Its intended function is limited to prioritising claims for human review. This limitation is deliberate and forms part of the system's ethical and regulatory design.

9. Conclusion

Digitalisation is changing insurance management by enabling policy information, customer services, claims processes, and analytical activities to be managed through integrated digital systems. Academic research demonstrates that machine-learning methods can assist with insurance fraud detection, including supervised classification, unsupervised anomaly detection, neural networks, and other predictive approaches.

However, insurance AI cannot be treated simply as a conventional classification problem. Insurance decisions can have significant consequences for individuals, and regulatory and supervisory frameworks increasingly emphasise fairness, transparency, explainability, data governance, record keeping, and human oversight. The EU AI Act explicitly identifies AI used for risk assessment and pricing in life and health insurance as a high-risk use case and establishes requirements for transparency, human oversight, and record keeping for high-risk AI systems

The proposed Insurance Management System therefore adopts a deliberately conservative AI architecture. Policy management, premium calculation, claims submission, payments, and claim tracking are handled through conventional software components, while machine learning is used

only to generate a fraud-risk flag or score for human review. The AI system does not independently approve, reject, or settle claims.

This design makes the project technically achievable for a BCA final-year student while remaining consistent with the direction of published research and regulatory thinking. The project demonstrates how AI can augment insurance operations without incorrectly presenting an experimental student model as an autonomous insurance decision-maker.

The principal contribution of the proposed system is therefore not the claim that AI can replace insurance professionals. Rather, it demonstrates a practical **human-in-the-loop model** in which digital insurance management is combined with AI-assisted fraud screening, while the final responsibility for consequential insurance decisions remains with an authorised human reviewer.

References

- [1] European Insurance and Occupational Pensions Authority (EIOPA), “The Trend Towards the Increasing Digitalisation of the Insurance Sector,” contribution to Eurofi Magazine, February 2024.
- [2] European Insurance and Occupational Pensions Authority (EIOPA), Consultative Expert Group on Digital Ethics in Insurance, “Artificial Intelligence Governance Principles: Towards Ethical and Trustworthy Artificial Intelligence in the European Insurance Sector,” Publications Office of the European Union, 2021. DOI: 10.2854/49874.
- [3] Jörn Debener, Volker Heinke, and Johannes Kriebel, “Detecting Insurance Fraud Using Supervised and Unsupervised Machine Learning,” **Journal of Risk and Insurance**, Vol. 90, No. 3, pp. 743–768, 2023. DOI: 10.1111/jori.12427.
- [4] Stijn Viaene, Guido Dedene, and Richard A. Derrig, “Auto Claim Fraud Detection Using Bayesian Learning Neural Networks,” **Expert Systems with Applications**, Vol. 29, No. 3, pp. 653–666, 2005. DOI: 10.1016/j.eswa.2005.04.030.
- [5] Mark Anthony Caruana and Liam Grech, “Automobile Insurance Fraud Detection,” *Communications in Statistics: Case Studies, Data Analysis and Applications*, Vol. 7, No. 4, pp. 520–535, 2021. DOI: 10.1080/23737484.2021.1986169.
- [6] Chamal Gomes, Zhuo Jin, and Hailiang Yang, “Insurance Fraud Detection with Unsupervised Deep Learning,” **Journal of Risk and Insurance**, Vol. 88, No. 3, pp. 591–624, 2021. DOI: 10.1111/jori.12359.
- [7] European Insurance and Occupational Pensions Authority (EIOPA), Consultative Expert Group on Digital Ethics in Insurance, “Artificial Intelligence Governance Principles: Towards Ethical and Trustworthy Artificial Intelligence in the European Insurance Sector,” Publications Office of the European Union, 2021. DOI: 10.2854/49874.
- [8] European Insurance and Occupational Pensions Authority (EIOPA), “Opinion on Artificial Intelligence Governance and Risk Management,” 6 August 2025.
- [9] European Insurance and Occupational Pensions Authority (EIOPA), “Digitalisation in Insurance: Keeping Track on Current Pace of Change,” contribution to **Eurofi Magazine**, April 2023.
- [10] European Insurance and Occupational Pensions Authority (EIOPA), “Consumer Trends Report 2022,” Publications Office of the European Union, 2022.
- [11] Faheem Aslam, Ahmed Imran Hunjra, Zied Ftiti, Wael Louhichi, and Tahira Shams, “Insurance Fraud Detection: Evidence from Artificial Intelligence and Machine Learning,” *Research in International Business and Finance**, Vol. 62, Article 101744, 2022. DOI: 10.1016/j.ribaf.2022.101744.