

Toward Integrated AI-Based Automation of Vehicle Insurance Claims: A Conceptual Framework and Preliminary Study

Sachin Rajput¹, Dr. Arshiya A. Khan²

Department of Computer Science and Applications, JSPM University, Pune, India

sachinrajput722446@gmail.com¹ ak.scos@jspmuni.ac.in²

Abstract

Motor vehicle insurance in India continues to rely on claim-handling procedures that were designed for a much smaller and slower-moving market: an adjuster inspects the vehicle, paperwork is gathered and checked by hand, and fraud is screened using rules of thumb built from experience rather than data. Individually, machine learning, computer vision, and natural language processing have each been shown capable of speeding up one piece of this workflow damage estimation, document reading, or fraud scoring yet the published work in this space almost always treats these three pieces as separate problems, leaving a person to stitch the outputs together and make the final call. In this paper, we present a preliminary, concept-stage contribution that precedes our full technical study. We examine the reasons claims automation research has remained fragmented, set out the requirements that a genuinely end-to-end system would need to satisfy, and lay out a conceptual architecture spanning intake, vision-based damage assessment, document understanding, fraud scoring, and a unifying decision layer that we intend to close this gap with. Rather than reporting experimental results, we establish the problem formulation, position our proposal against existing literature, and outline the methodology, datasets, and evaluation plan we will use in our subsequent full-length study, in which we build and evaluate the system.

Keywords: *insurance technology, claims automation, machine learning, computer vision, natural language processing, fraud detection, decision support, motor insurance, artificial intelligence, insurtech.*

I. INTRODUCTION

India's vehicle parc has expanded rapidly over the last decade, and the motor insurance sector has grown alongside it, which means the volume of claims an insurer must handle every year has grown too [1], [2]. The underlying claims-processing workflow has not kept pace, the underlying process by which a claim moves from a first notification of loss to a final settlement. In most insurers, that process still runs

through a sequence of manual checkpoints: a surveyor or adjuster is dispatched or asked to review photographs, paperwork such as the policy, a First Information Report, and repair estimates is collected and cross-checked by a claims officer, and a fraud referral is raised only when something looks unusual to that reviewer [3], [4]. Every one of those checkpoints depends on an individual's availability, workload, and judgment on a given day, resulting in substantial variation in turnaround time and consistency vary so much across otherwise similar claims [5], [4].

Fraud compounds the problem. Even though fraudulent claims are widely believed to represent a modest share of total claim volume, the losses attributable to them are large in absolute terms simply because of the large number of claims processed each year through the system each year [6], [7]. Detecting these cases by manual review alone is difficult, since a reviewer working through a high volume of claims has little time to notice a subtle inconsistency, such as a repair date that does not line up with the accident date on record [7].

Artificial intelligence has already demonstrated, in isolated studies, that it can outperform manual methods on individual pieces of this workflow: convolutional and transformer-based vision models can localize and grade vehicle damage from photographs [5], [8], [9], optical character recognition combined with modern language models can pull structured facts out of scanned paperwork [10], and ensemble learning methods can flag anomalous claims more reliably than static rule sets [6], [11]. The difficulty is that these capabilities have mostly been developed and evaluated as standalone systems [12], [13], [14]. A damage-grading model does not talk to a fraud-scoring model, and neither is wired into a document-verification pipeline, so a person still has to sit in the middle and combine their outputs before a claim can be resolved.

In this paper, we step back from building and evaluating such a system that is the subject of our companion full-length paper and instead do the groundwork that ought

to precede it: we analyze the factors contributing to this fragmentation, define the functional requirements of an integrated system, and propose a conceptual architecture for it. Our intent is to make the case for this design before we commit to the implementation details and empirical claims that will follow.

The subsequent sections are organized as follows. In Section II, we motivate the problem in more depth. In Section III, we survey the closest prior work and explain the limitations preventing a fully integrated solution of an integrated solution. In Section IV, we formalize the research gap. In Section V, we introduce our proposed conceptual framework. In Section VI, we outline our planned methodology and evaluation approach for the subsequent study. In Section VII, we discuss the anticipated contributions; in Section VIII, we note the scope and limitations of this preliminary work; and in Section IX, we conclude.

II. BACKGROUND AND MOTIVATION

Understanding the factors limiting claims-processing automation has resisted automation is to examine the processing requirements of a typical claim. First, someone has to determine the extent of vehicle damage and what it will cost to repair traditionally done by a surveyor's visual inspection [5], [3]. Second, someone has to confirm that the paperwork supporting the claim is genuine and internally consistent: does the registration number on the policy match the one on the FIR, does the repair invoice correspond to the damage described, is the accident date plausible given when the FIR was filed [10], [3]. Third, one has to judge whether the overall pattern of the claim looks like the many other legitimate claims the insurer has settled before, or whether it resembles patterns associated with staged accidents, inflated repair bills, or duplicated documentation [6], [7]. Only after all three judgments are made can a decision to approve, deny, or escalate be reached [3], [4].

Each of these three judgments is, in principle, a task that a trained model can perform: damage grading is an image classification and localization problem, document verification is an information-extraction and consistency-checking problem, and fraud screening is a pattern-recognition problem over structured claim histories. The relationships among these tasks are not immediately apparent from any one task in isolation, however the fraud signal often depends on the document-verification outcome (for example, a mismatched date is itself a fraud indicator), and the final

decision has to weigh all three judgments against one another rather than treating them as independent votes. We treat this interdependence as the central design challenge that motivates the integrated, rather than modular but disconnected, system we propose in Section V.

III. RELATED WORK

Research directly addressing automated fraud detection in insurance has shown that predictive and anomaly-based models can substantially cut both processing time and missed-fraud rates when applied narrowly to claims data [6], but such studies typically stop at the fraud-scoring boundary and do not incorporate damage assessment or document handling into the same pipeline.

A separate stream of work has examined the automation of vehicle and property damage estimation using convolutional architectures [5], generally finding that automated grading is both faster and more repeatable than manual surveying, though again as a self-contained capability rather than one component of a larger decision system. Related studies on cloud-native deployment of AI for claims processing [15] have argued convincingly that scalable infrastructure is a precondition for any production system, without proposing a vehicle-specific or fully integrated pipeline themselves.

Santhoshkumar's work focused on motor and floater policies has explored robotic process automation and conversational interfaces to improve the customer-facing side of claims handling [16], but without a computer-vision component for damage assessment. The study that comes closest to an integrated design combines machine learning, computer vision and natural language processing on top of an existing commercial claims platform [13], reporting a meaningful reduction in manual touchpoints; its reliance on that platform's legacy architecture, however, has been reported to create integration friction, and performance is noted to be sensitive to inconsistent real-world lighting conditions in submitted photographs [13].

Broader survey papers on AI adoption in motor insurance corroborate that automation can cut processing time substantially at an industry level [14], but these are reviews of the landscape rather than working implementations. Other studies on commercial auto insurance make a case for big-data and predictive-

analytics approaches to risk and fraud without specifying an implementable architecture [17], while proof-of-concept damage-detection systems built on modern object-detection frameworks have reported usable accuracy without incorporating any fraud-screening capability at all [18], and enterprise-integration studies have examined how predictive analytics and language processing can be layered onto existing claims platforms without adding new automated decision-making capability [19], [12].

Taken together, this body of work establishes two things clearly: first, that each of the three underlying tasks damage grading, document verification, and fraud scoring is individually tractable with current AI methods [6], [5], [10], and second, that no published study has combined all three into a single system in which one module's output informs another can inform another's and where the final approve, deny, or escalate decision is itself made by the system rather than handed back to a adjuster at the end of the pipeline [12]-[14], [18], [19]. This is the gap that motivates the framework we propose in Section V.

IV. RESEARCH GAP AND PROBLEM FORMULATION

Based on the review above, four specific gaps can be identified:

First, almost every existing study addresses damage assessment, document processing, or fraud detection in isolation, rather than as parts of one connected pipeline [6], [5], [16]. Second, in cases where automation exists across more than one of these tasks, it is typically bolted onto legacy claims-management infrastructure rather than designed from the ground up to pass information between modules [13]. Third, none of the systems reviewed make the final claim decision autonomously; an adjuster remains the last step even when every upstream task has already been automated [18], [19]. Fourth, published throughput figures for automated components are rarely translated into an end-to-end turnaround-time comparison against the days-long manual baseline, making it hard to judge the practical benefit of adopting any single piece in isolation [14], [4].

We argue that the primary objective of this research should extend beyond improving the three individual tasks, but designing a system in which damage assessment, document verification, and fraud scoring share a common data flow and feed a single automated decision layer capable of resolving straightforward

claims without a person's involvement, while still routing ambiguous cases to a reviewer. This is the design goal we take up next.

V. PROPOSED CONCEPTUAL FRAMEWORK

To address the gap above, we propose a five-stage architecture coordinated through a shared orchestration layer, so that the output of any one stage is available as an input to the stages that follow rather than being consumed only by a reader.

A. Claim Intake and Pre-processing

We design the entry point of the system to capture claim details, vehicle photographs, and supporting documents through a claimant-facing interface, validate that the minimum required information has been supplied, and normalize it into a structured record that downstream modules can consume.

B. Vision-Based Damage Assessment

We plan to process submitted photographs through an object-detection stage to localize damaged regions [8], followed by a classification stage to grade the type and severity of damage using a residual network architecture [9], with an associated cost estimate derived from that grading [5]. We choose an object-detector-plus-classifier pairing for this stage because modern detectors remain robust to the cluttered, inconsistently lit images that claimants typically submit from a phone camera [18].

C. Document Understanding

We plan to convert scanned policies, First Information Reports, and repair invoices to text via optical character recognition and then parse them with a named-entity-recognition model fine-tuned on insurance documents [10], so that fields such as policyholder identity, vehicle registration, incident date, and claimed amount can be extracted automatically. We also plan a validation step to cross-check these extracted fields against the policy database and to flag internal inconsistencies, such as a First Information Report dated implausibly relative to the reported accident date [3].

D. Fraud Risk Scoring

We intend to combine claim-history features, the outputs of the document-consistency checks, and derived damage-cost signals into a fraud-risk score [6], [11]. Because fraud indicators are comparatively rare relative to legitimate claims, we plan an ensemble approach combining complementary tree-based learners [11], [20], on the premise that different model families

tend to surface different feature interactions in imbalanced data [20].

E. Automated Decision Layer

In the final stage, we plan to combine the damage assessment, document-validity outcome, and fraud-risk score into a single weighted decision, resolving clearly favorable and clearly unfavorable claims automatically while routing ambiguous cases to an adjuster, along with an explanation of which factors most influenced the recommendation.

We deliberately describe this architecture at the responsibilities assigned to each stage and the expected information flow between stages between stages, rather than at the level of trained model weights or measured accuracy we reserve that level of detail, together with the datasets, training procedure, and quantitative evaluation, for our full-length companion paper.

VI. PLANNED METHODOLOGY AND EVALUATION

In our subsequent full study, we plan to evaluate the framework using a labeled image dataset for damage grading [5], [9], a historical claims dataset for fraud-model training [6], [20], and a corpus of scanned insurance documents for the document-understanding component [10], with a test split, training and validation appropriate to each. We plan to evaluate the system using standard classification metrics like accuracy, precision, recall, F1-score, and area under the ROC curve computed per module, together with a system-level measure of end-to-end processing time benchmarked against the conventional multi-day manual baseline [14], [4]. We also plan ablation comparisons against single-model baselines for the fraud-scoring component, to establish whether an ensemble design meaningfully outperforms its constituent models rather than assuming it by default [11], [20].

VII. ANTICIPATED CONTRIBUTIONS

If realized as described, we expect our framework to contribute in four ways: a claims pipeline in which we design damage assessment, document verification, and fraud scoring to exchange information rather than operate independently; an automated decision layer that can resolve a meaningful share of claims while appropriately flagging ambiguous cases for further processing; an architecture we evaluate explicitly against an end-to-end turnaround-time baseline rather

than only against per-module accuracy figures; and a design we tailor specifically for the Indian motor-insurance context, including compliance and data-handling considerations relevant to that setting.

VIII. SCOPE AND LIMITATIONS OF THIS STUDY

We deliberately scope this paper as a preliminary contribution. We do not report trained-model performance, dataset statistics, or system benchmarks here, since all of that belongs to the empirical study that follows it. Our purpose in this paper is instead to establish the problem formulation, justify our architectural choices against the existing literature, and set out our evaluation plan, so that the results we report in our subsequent paper can be read against a clearly stated set of expectations rather than being the first articulation of the problem.

IX. CONCLUSION

Manual vehicle-insurance claims processing in India remains slow, inconsistent, and vulnerable to fraud that is difficult for individual reviewers to catch. While machine learning, computer vision and natural language processing have each been shown to help with one piece of this workflow, the literature has not yet converged on a design in which these components are integrated into a single automated decision-making system. In this paper, we made the case for such a system: we reviewed the closest prior work to establish where it falls short of full integration, and we proposed a five-stage conceptual architecture intake, vision-based damage assessment, document understanding, fraud scoring, and an automated decision layer intended to close that gap. As the next stage of this research, we will implement and empirically evaluate this architecture in a full-length a design in which these components are integrated and measure the modules described here against the manual baseline they are intended to improve upon.

REFERENCES

- [1] Ministry of Road Transport and Highways (MoRTH), Road Transport Year Book 2022–2023, Government of India.
- [2] Insurance Regulatory and Development Authority of India (IRDAI), Annual Report 2023–2024, Hyderabad, India.
- [3] Insurance Information Institute (III), Auto Insurance Claims Process, 2023.

- [4] McKinsey & Company, Insurance 2030 — The Impact of AI on Claims Management, 2021.
- [5] K. Rahul, "Automated Vehicle and Property Loss Assessment Using Convolutional Neural Networks," IEEE Trans. Comput. Intell. AI Games, vol. 12, no. 3, pp. 88–99, 2020.
- [6] M. Shaik, "Intelligent Automation in Insurance: Machine Learning Models for Fraud Detection and Predictive Analytics," J. Insurance Technol., vol. 7, no. 2, pp. 45–62, 2019.
- [7] Coalition Against Insurance Fraud, The Cost of Insurance Fraud, 2023.
- [8] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," arXiv:1804.02767, 2018.
- [9] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. IEEE CVPR, 2016, pp. 770–778.
- [10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in Proc. NAACL-HLT, 2019, pp. 4171–4186.
- [11] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in Proc. 22nd ACM SIGKDD, 2016, pp. 785–794.
- [12] V. Komperla, "AI-Enhanced Claims Processing: Business Models and Operational Efficiency in Insurance Organizations," Int. J. Insurtech, vol. 3, no. 1, pp. 12–28, 2021.
- [13] S. Gollapudi, "End-to-End AI Framework for Insurance Claims: Integrating ML, NLP, and Computer Vision with Guidewire," J. Enterprise AI Syst., vol. 2, no. 1, pp. 17–33, 2025.
- [14] L. Sun et al., "A Comprehensive Review of AI Applications in Insurance Claims Management," AI Finance Insur. Surv. J., vol. 6, no. 2, pp. 109–131, 2025.
- [15] R. Devaraj, "Cloud-Based AI Integration for Claims Processing Automation in Property and Casualty Insurance," J. Cloud Comput. Insur. Appl., vol. 5, no. 4, pp. 201–218, 2023.
- [16] P. Santhoshkumar, "AI Applications in Floater Motor Insurance: RPA, Chatbots, and ML for Claims Handling," Insur. Technol. Rev., vol. 9, no. 1, pp. 34–52, 2024.
- [17] H. Ma, "AI in Commercial Auto Insurance: Risk Assessment, Fraud Detection, and Customer Service Enhancement," J. Actuarial Sci. AI, vol. 11, no. 3, pp. 77–94, 2024.
- [18] C. Jayalath, "Smart Claim AI: Deep Learning for Vehicle Damage Detection and Cost Estimation," in Proc. Int. Conf. Appl. Mach. Learn., 2023, pp. 312–319.
- [19] K. Rahul, "AI-Based Insurance Platforms: Predictive Analytics, NLP, and Real-Time Processing for Claims Management," Enterprise Syst. J., vol. 14, no. 2, pp. 55–68, 2022.
- [20] N. V. Chawla et al., "SMOTE: Synthetic Minority Over-sampling Technique," J. Artif. Intell. Res., vol. 16, pp. 321–357, 2002.