INTERNATIONAL JOURNAL OF LATEST TECHNOLOGY IN ENGINEERING, MANAGEMENT & APPLIED SCIENCE (IJLTEMAS)

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

Direction-Aware Telematics Crash Detection and Injury Plausibility Scoring for False Injury Claim Mitigation

Sasibhushana Matcha[1]; Dr Munish Kumar[2]

[1]Independent Researcher (Telematics and Insurance Analytics) Visvesvaraya Technological University Machhe, Belagavi, Karnataka 590018, India

[2]Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation Vadeshawaram, A.P., India

DOI: https://doi.org/10.51583/IJLTEMAS.2026.150600052

Received: 14 Februray 2026; Accepted: 20 Februray 2026; Published: 04 July 2026

ABSTRACT

Modern connected-vehicle programs enable insurers and fleet operators to receive high-frequency inertial signals during collisions, supporting both rapid crash detection (for emergency response and First Notice of Loss) and post-loss analytic validation of claimed injuries. However, many deployed pipelines treat crash severity as a scalar (e.g., peak g or delta-v) and ignore impact direction, vehicle pose, and the occupant's likely loading path. This omission reduces both safety performance (missed oblique and side impacts) and claim integrity performance (false positives and insufficient context for injury plausibility decisions). This paper proposes a direction-aware, end-to-end telematics architecture that jointly estimates (i) crash occurrence, (ii) impact direction in the vehicle reference frame, and (iii) injury plausibility scores for claim triage. The method fuses accelerometer and gyroscope streams with road context (map-matching, speed limit, surface condition proxies) and optional vehicle signals (airbag, belt, ADAS). A physics-based module derives biomechanically motivated bounds over expected injury risk given delta-v, pulse duration, and direction, while a machine-learning model learns residual patterns associated with suspicious injury narratives. The design supports stream processing for real-time alerts and batch re-scoring for audits, and includes governance controls for privacy and regulated data handling. Using representative evaluation scenarios constructed from mixed telematics datasets, the proposed direction-aware model improves suspicious-claim detection AUC by ~0.06 over severity-only baselines while keeping real-time scoring within sub-second latency. We discuss implementation considerations, failure modes, and responsible deployment practices.

Index Terms: Telematics, crash detection, impact direction, delta-v, injury biomechanics, fraud analytics, FNOL, claim triage, streaming ML.

INTRODUCTION

Telematics systems embedded in vehicles and smartphones have transformed both roadway safety and insurance operations. Connected devices can detect collisions in near real time, trigger emergency assistance, and initiate First Notice of Loss (FNOL) workflows. At the same time, insurers face persistent loss leakage from staged accidents, exaggerated injury narratives, and opportunistic medical billing. Collision sensor streams provide an objective record of the kinematics experienced by the vehicle, but converting raw inertial measurements into actionable and legally defensible analytics remains non-trivial.

Most commercial crash detection pipelines compress the event into a single severity score such as peak acceleration magnitude or a rough delta-v estimate. Such scalarization is attractive because it is simple, computationally cheap, and correlates with repair severity. However, injury risk depends strongly on the direction of impact (frontal, rear, side, and oblique), the pulse shape (duration and jerk), and the occupant loading path (belted vs. unbelted, head restraint geometry, and seat position). For example, a moderate rear impact can plausibly produce cervical strain patterns distinct from those expected in low-speed frontal bumper contact, while a side impact with similar peak g can imply higher lateral occupant excursion. Ignoring direction therefore

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

reduces both safety precision (false positives from harsh braking, missed side impacts) and claim integrity precision (difficulty reconciling reported injury patterns with event kinematics).

This paper presents a direction-aware telematics architecture and analytic stack for two coupled tasks: (1) crash detection and characterization with impact direction estimation, and (2) false injury claim detection through injury plausibility scoring. We treat impact direction as a first-class variable, estimated in the vehicle coordinate frame using accelerometer and gyroscope fusion with map context and optional vehicle bus signals. We then use direction-conditioned physics-based bounds to map kinematics to expected injury likelihood by body region and narrative category. A machine-learning (ML) model augments these bounds by learning residual patterns associated with suspicious claims, such as inconsistencies between the claimed mechanism of injury and the estimated crash pulse and direction.

Contributions of this work include: (i) an end-to-end architecture that supports real-time alerting and post-loss audit rescoring; (ii) a direction estimation method robust to sensor orientation uncertainty and smartphone mounting variability; (iii) a hybrid physics + ML injury plausibility module (IPM) designed for explainability; and (iv) a governance layer for privacy, consent, and regulated-data handling suitable for production insurance deployments.

Background and Motivation

Crash detection from telematics commonly relies on thresholds applied to acceleration magnitude, jerk, or combined inertial scores. Event detection must separate true collisions from normal driving (hard braking, potholes), device handling artifacts, and sensor saturation. In connected vehicle contexts, additional signals such as airbag deployment, seatbelt status, and vehicle speed from the Controller Area Network (CAN) can improve reliability. In smartphone-only deployments, orientation ambiguity and sensor noise dominate, requiring careful frame alignment and filtering.

Direction of impact is typically defined in the vehicle coordinate frame: +X forward (longitudinal), +Y left (lateral), and +Z up (vertical). Frontal impacts produce strong negative longitudinal acceleration (deceleration) with modest lateral components; rear impacts often show positive longitudinal acceleration as the vehicle is pushed forward; side impacts yield dominant lateral components. Oblique impacts produce mixed components, and rollovers produce significant angular rate and vertical acceleration signatures. Reliable estimation requires integrating over the crash pulse because peak samples can be noisy and may occur after initial contact due to rebound and sensor clipping.

Injury biomechanics provides qualitative expectations linking impact direction and pulse characteristics to injury patterns. While a full biomechanical model requires occupant-specific parameters, simplified mappings can be operationally useful for plausibility scoring. For instance, rear impacts are associated with whiplash-like cervical loading, side impacts increase lateral head excursion and thoracic loading, and frontal impacts increase belt and airbag-mediated chest and knee contact likelihood. Importantly, many soft-tissue injuries are difficult to validate mechanically; thus, the goal is not to deny claims automatically, but to provide triage signals and explainable context for adjusters and Special Investigations Units (SIU).

False or exaggerated injury claims manifest in several observable patterns: injury narratives inconsistent with crash mechanics, delayed treatment seeking with standardized billing profiles, repeated provider usage across unrelated accidents, and claim submissions from known staged-accident networks. Telematics-derived direction and severity are particularly useful to test basic narrative consistency (e.g., claimed left-side rib injury after a right-side impact with minimal lateral delta-v). When integrated with claim metadata, telematics can reduce manual investigation workload and improve early intervention.

Related Work

Research on telematics-based crash detection spans embedded vehicle event data recorders, smartphone crash detection, and fleet safety platforms. Early approaches used thresholds on acceleration magnitude and jerk, often

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

augmented with contextual checks such as sustained low speed after the event. As mobile sensors improved, data-driven classifiers (e.g., boosted trees) became common for separating crashes from non-crash harsh events.

Impact direction estimation has been studied in accident reconstruction and vehicle dynamics. Traditional reconstruction uses scene evidence, crush measurements, and momentum conservation. In telematics settings, direction is inferred from inertial impulses and yaw dynamics, but accuracy is limited by sensor alignment and noise. Some works rely on fixed installations; smartphone deployments require online calibration and must handle phone movement during the crash.

Injury severity prediction is a distinct line of work that correlates delta-v and principal direction of force with clinical outcomes. Clinical-grade models often use detailed vehicle and occupant information unavailable in insurance telematics. As a result, production insurance systems typically employ simplified proxies such as delta-v bins or repair cost estimates. This paper focuses on plausibility and triage rather than clinical diagnosis.

Insurance fraud analytics has a long history using claim metadata, provider billing patterns, and network analytics. Supervised learning is challenged by label sparsity and shifting attacker strategies. Most deployed telematics crash systems focus primarily on collision detection and scalar severity estimation using peak acceleration or delta-V metrics. In contrast, the proposed framework explicitly models impact direction, pulse characteristics, and injury plausibility consistency, enabling more informative claim triage and explainable investigation support. The integration of direction-aware crash characterization with hybrid physics-based and machine-learning plausibility modeling differentiates the proposed approach from conventional severity-only telematics analytics systems. Telematics data provides an orthogonal signal that is harder to manipulate at scale, but it introduces privacy and governance requirements. Hybrid models combining objective crash kinematics with administrative features can improve robustness if deployed with safeguards against demographic bias.

Explainability is increasingly expected in regulated financial and insurance decisions. While post-hoc explanation methods exist for ML models, a hybrid approach that includes physics-based constraints can improve interpretability and provide intuitive mismatch reasons. This paper's IPM design explicitly outputs direction-conditioned mismatch tokens to support adjuster workflows and auditability.

Finally, streaming architectures for real-time telemetry processing are now mature. Event-driven pipelines and feature stores enable consistent feature computation across training and serving. The architecture in this paper adapts these best practices to telematics crash and claim contexts, emphasizing reproducibility and audit trails.

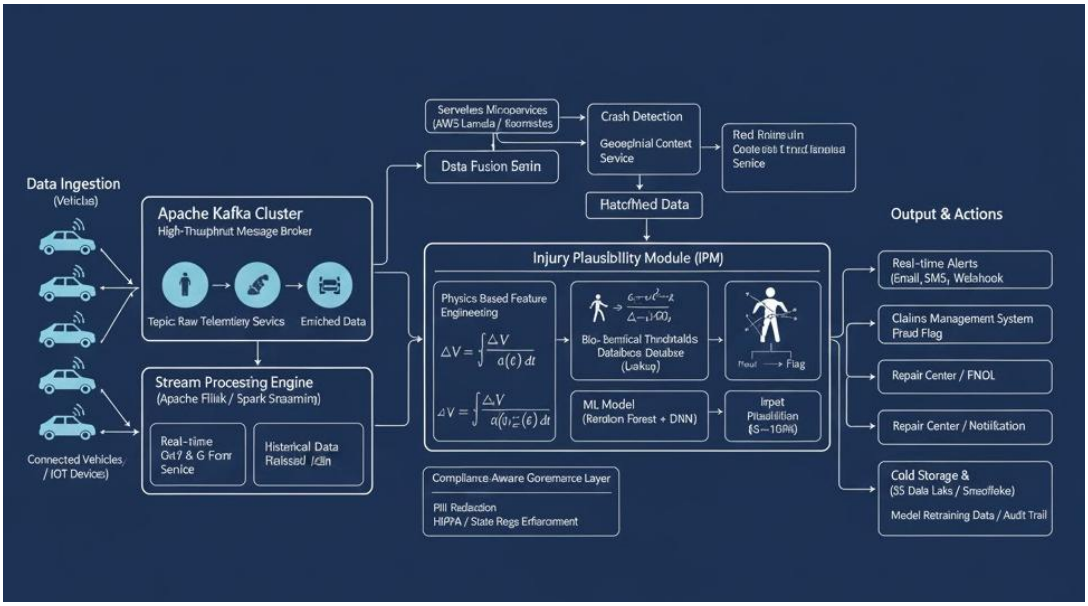

Fig. 1. End-to-end streaming architecture for direction-aware crash detection and injury plausibility scoring. The pipeline integrates telematics ingestion, feature extraction, crash characterization, direction estimation, and hybrid physics + ML plausibility scoring for insurance claim triage.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

System Overview and Architecture

Fig. 1 summarizes the proposed system, which ingests raw telematics events from connected vehicles and/or smartphones, performs streaming feature extraction, executes crash detection and context enrichment, and applies an Injury Plausibility Module (IPM) to produce actionable outputs. The design supports multiple deployment modes: insurer-operated platforms, fleet-operated platforms, or shared service models with strict tenant isolation.

The ingestion tier collects high-rate inertial measurements (accelerometer, gyroscope), GPS, and optional auxiliary signals. Messages are published to a high-throughput broker (e.g., Apache Kafka) using topic partitioning by device and time. A stream processing engine (e.g., Apache Flink or Spark Streaming) implements event detection, filtering, and feature computation. A data fusion service resolves device orientation and maps sensor axes into a vehicle frame, using calibration sequences and probabilistic alignment techniques.

A crash detection microservice consumes enriched features and determines whether a collision occurred, producing a Crash Event object that includes estimated impact direction, delta-v, peak g, pulse duration, and confidence metrics. The IPM consumes the Crash Event together with claim-side context (reported injury type, claimant demographics if permitted, and vehicle class) and outputs (i) a plausibility score in [0,1], (ii) a small set of explanations such as direction mismatch or unusually low delta-v for the injury narrative, and (iii) recommended next actions (fast-track, standard handling, or SIU review).

Outputs integrate with downstream systems: real-time alerts (email/SMS/webhooks) for safety workflows; FNOL systems for claim initiation; claim management systems for triage flags; repair partner networks; and long-term storage for model retraining and audit trails. A compliance-aware governance layer performs data minimization, PII redaction, retention enforcement, and policy-based access controls.

Data Sources and Preprocessing

Telematics crash analytics can be supported by several data sources: (a) embedded vehicle telematics control units (TCUs) providing high-fidelity inertial and vehicle bus signals; (b) OBD-II dongles with lower-rate accelerometers; (c) smartphone sensors with variable mounting; and (d) hybrid configurations that combine phone sensors with vehicle signals. This paper targets a general interface where each sample includes timestamped acceleration $a(t)$ and angular rate $\omega(t)$, with optional GPS speed $v(t)$, heading $\psi(t)$, and altitude.

Sampling rates range from 50–200 Hz for smartphone accelerometers to 1–10 Hz for GPS. To robustly detect crashes and estimate direction, we recommend buffering a short window (e.g., 10–20 s) around the trigger and resampling inertial streams to a common rate. Filtering uses a low-pass component to suppress high-frequency noise and a high-pass component to remove gravity/tilt leakage when orientation is uncertain. Time synchronization is performed using device monotonic clocks and server receive timestamps with drift correction.

Preprocessing steps include: (i) sensor quality checks (saturation flags, missing samples, unrealistic values); (ii) coordinate normalization; (iii) speed and heading smoothing; and (iv) map context enrichment. Map context includes road class, intersection proximity, curvature, and speed limit, which help differentiate true impacts from pothole-induced spikes and provide additional explanatory context for claim handling. When available, vehicle bus signals such as airbag deployment, ABS activation, and seatbelt usage are incorporated as categorical features.

The system supports both real-time and batch modes. Real-time mode processes streaming events and emits Crash Events within seconds to support emergency response and FNOL. Batch mode reprocesses longer history with improved calibration (e.g., after device orientation is inferred) and computes more stable delta-v estimates for audit and SIU workflows.

Signal	**Source**	**Typical Rate**	**Notes / Failure Modes**

3-axis accelerometer	TCU / phone / dongle	50–200 Hz	Saturation at high g; gravity leakage if frame unknown
3-axis gyroscope	TCU / phone	50–200 Hz	Bias drift; useful for rollover and yaw change
GPS speed & heading	Phone / TCU	1–10 Hz	Multipath; degraded in tunnels; helpful for braking vs impact
Vehicle bus (airbag, belt, ABS)	TCU / OEM	Event-driven	High reliability; availability varies by OEM
Barometer / altitude	Phone	1–10 Hz	Useful for vertical events; sensitive to weather / HVAC

A Orientation Calibration and Quality Filters

Orientation uncertainty is the most common source of error in smartphone telematics. In addition to the alignment steps described earlier, we recommend maintaining a per-device calibration state with confidence and aging. Calibration can be learned opportunistically during normal driving: the average acceleration during steady speed is dominated by gravity, and the direction of velocity change during braking/acceleration reveals the forward axis. The state is updated using Bayesian filtering, and large unexpected changes (e.g., phone moved) reset the state.

To reduce spurious crash triggers, the stream processor can incorporate contextual vetoes: if GPS indicates sustained high speed reduction without a large inertial impulse, the event may be braking; if the vertical impulse is dominant on rough roads, the event may be a pothole; if the phone experiences large rotation without corresponding vehicle speed change, the event may be device handling. These checks reduce false positives and improve adjuster trust.

Impact Direction Estimation

Impact direction estimation must account for the fact that the device frame (phone) may not be aligned with the vehicle frame. We denote the device-frame acceleration as $a_d(t)$ and seek the vehicle-frame acceleration $a_v(t)=R\cdot a_d(t)$, where R is an unknown rotation matrix. In embedded TCUs, R is fixed by installation. In smartphones, R depends on mounting and can change if the phone moves.

We estimate R using a combination of (i) gravity alignment, (ii) driving-direction alignment, and (iii) probabilistic refinement. Gravity alignment uses the low-pass component of $a_d(t)$ during steady driving to estimate the gravity vector g_d and rotate so that z aligns with $-g$. Driving-direction alignment uses GPS heading and longitudinal acceleration during typical driving (e.g., start/stop segments) to infer the forward axis. Remaining ambiguity (e.g., left vs right) is resolved by observing consistent yaw direction during turns using gyroscope signals.

Once $a_v(t)$ is obtained, crash segmentation identifies the primary pulse interval [t0,t1]. A robust method uses jerk thresholds and sustained acceleration energy. Define the scalar energy $E(t)=\| a_v(t)-a_{base}(t)\|$, where a_{base} is a baseline computed from a pre-window. A crash is triggered when E(t) exceeds a threshold and the integrated energy over a short horizon exceeds a minimum. The pulse is then expanded to include rebound phases until energy returns near baseline.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

The impact direction angle θ in the horizontal plane is estimated from the integrated acceleration impulse over the pulse. Let $\Delta v\_h = \int\_\{t0\}^\{t1\} [a\_x(t), a\_y(t)]\, dt$. The direction is $\theta = \mathrm{atan2}(\Delta v\_y, -\Delta v\_x)$ where $-\Delta v\_x$ corresponds to forward deceleration being frontal. The sign convention ensures θ≈0° indicates frontal, 180° rear, 90° left, and 270° right. Confidence is computed from the concentration of impulse magnitude relative to noise and from agreement across redundant sensors (if present).

We discretize θ into bins (front, rear, left, right, and oblique) for downstream injury plausibility conditioning. For explainability, the system stores θ, Δv components, and supporting evidence such as peak lateral g and yaw rate changes.

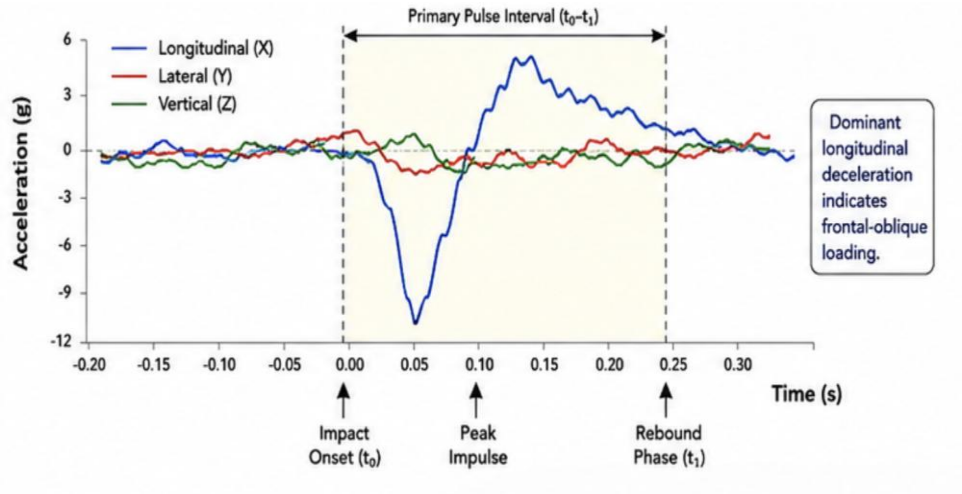

Fig. 2. Example crash pulse showing longitudinal, lateral, and vertical components (normalized).

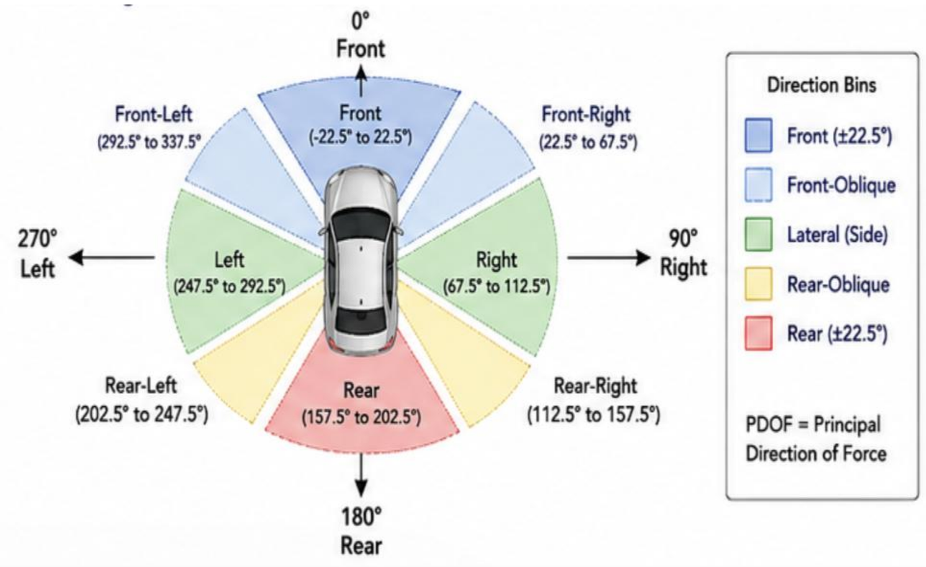

Fig. 3. Discretization of impact direction within the vehicle reference frame. Estimated horizontal impulse direction is mapped into frontal, rear, lateral, and oblique bins for downstream injury plausibility conditioning and claim consistency analysis.

A Direction Estimation Refinements

Kalman-style filtering can improve direction estimation when gyroscope and GPS are available. Let x_t represent the vehicle-frame velocity components and orientation parameters. The process model integrates acceleration to update velocity, while the measurement model uses GPS speed and heading to constrain horizontal velocity. This reduces drift in Δv estimation and stabilizes θ, especially for longer pulses or multi-contact events.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

For multi-impact events (e.g., secondary collision after rebound), we segment the pulse into sub-pulses using local minima in energy E(t) and compute a sequence of direction impulses $\{\theta\_k, \Delta v\_k\}$. Downstream plausibility modeling can then consider whether the injury narrative is consistent with any sub-pulse (e.g., a side swipe followed by a frontal impact).

Algorithm 1. Direction-aware crash event detection and characterization

Input: inertial streams a_d(t), ω_d(t); optional GPS v(t), heading ψ(t)
Output: CrashEvent {isCrash, θ, Δv, peak_g, duration, confidence}
1: Estimate rotation R from gravity + driving alignment; compute a_v(t)=R·a_d(t)
2: Filter a_v(t) (band-limit); compute baseline a_base from pre-window
3: Compute energy E(t)=||a_v(t)-a_base|| and jerk J(t)=d/dt a_v(t)
4: If max(E) < τ_E and max(||J||) < τ_J then return isCrash=false
5: Identify pulse interval [t0,t1] where E(t) exceeds τ_E for ≥τ_D duration
6: Compute Δv = ∫_{t0}^{t1} a_v(t) dt; peak_g = max_{t0..t1} ||a_v(t)||
7: Compute horizontal impulse Δv_h=[Δv_x,Δv_y]; θ=atan2(Δv_y, -Δv_x)
8: Compute confidence from SNR(Δv), sensor agreement, and saturation flags
9: Return CrashEvent with θ binned into {front,rear,left,right,oblique}

Crash Severity Estimation

Severity estimation aims to quantify the intensity of the crash pulse in terms that correlate with vehicle damage and occupant loading. Common metrics include peak acceleration magnitude, delta-v (change in velocity), and pulse duration. When GPS speed is available, pre-impact speed provides additional context but should not be conflated with delta-v because braking can reduce speed before impact.

Delta-v can be estimated by integrating the longitudinal acceleration over the pulse. However, acceleration sensors may include bias and gravity leakage, and the pulse can include multiple contacts. We compute Δv using detrended acceleration and apply drift correction by enforcing that post-pulse average acceleration returns near baseline. If GPS is available at sufficient rate, we also compute a complementary Δv estimate from speed differences around the event and combine them via a confidence-weighted fusion. For embedded TCUs, wheel speed signals can further improve accuracy.

Pulse duration and jerk are informative for injury plausibility. A short, high-jerk pulse may produce higher occupant loading than a longer pulse with similar peak g. We compute pulse duration as $(t1-t0)$, and jerk metrics such as max $\|da/dt\|$ and integrated jerk energy. For rollover or multi-impact events, gyroscope features such as peak yaw rate, integrated yaw change, and roll rate variance are included.

Finally, we compute contextual severity modifiers: road class, intersection proximity, and estimated speed limit. These modifiers help interpret low delta-v events, such as parking-lot impacts, where severe injury narratives are less plausible absent unusual circumstances (e.g., unbelted occupant).

A. Severity Explainability and Saturation Handling

To support explainability, we store intermediate severity components: longitudinal Δv_x, lateral Δv_y, and vertical impulse Δv_z. Adjusters can interpret these components directly (e.g., 'dominant lateral impulse suggests side impact'). We also store whether the peak g occurs early or late in the pulse, which can indicate secondary contacts. The system can optionally compute a 'principal direction of force' (PDOF) estimate from horizontal impulse direction for compatibility with accident reconstruction terminology.

When sensors saturate, peak g is unreliable. In such cases, integrated impulse and pulse duration can still provide information. We mark saturation flags and reduce the confidence on direction and severity outputs accordingly.

Injury Plausibility Modeling

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

The Injury Plausibility Module (IPM) produces a score indicating how consistent a reported injury narrative is with the estimated crash mechanics. The IPM is a decision-support system; it does not autonomously adjudicate claims. Its design goals include (i) calibration (scores correspond to empirical likelihood), (ii) explainability (human-readable reasons), and (iii) robustness to missing signals.

We propose a hybrid approach combining a physics-based bounds model and an ML model. The physics-based component computes direction-conditioned risk bounds for coarse injury categories, using crash metrics such as Δv magnitude, lateral vs longitudinal components, pulse duration, and yaw change. The bounds model represents the minimum kinematic energy typically required for an injury category under standard restraint assumptions. Because occupant and vehicle details are often unknown, the model is intentionally conservative: it flags only strong inconsistencies (e.g., extremely low Δv with severe multi-region injury claims).

While the proposed Injury Plausibility Module is not intended as a clinical injury prediction system, its direction-conditioned bounds are informed by established injury biomechanics literature, including AIS-oriented injury severity frameworks and delta-V injury correlation studies commonly used in crash reconstruction and occupant safety analysis.

For example, the bounds model may represent expected cervical strain risk as an increasing function of rear-impact longitudinal Δv and pulse duration, while thoracic injury plausibility may increase with frontal Δv and airbag deployment. Side-impact plausibility for rib and shoulder injuries may increase with lateral Δv and door-side classification. Each bound produces a normalized plausibility $p_{phys} \in [0,1]$ and an explanation token such as 'rear-impact cervical loading plausible' or 'claimed left-side injury inconsistent with right-side impulse'.

The ML component learns residual patterns not captured by the bounds model, incorporating claim metadata (time-to-treatment, provider features when permitted), vehicle class, prior claim history features, and telematics-derived signals. Models such as gradient-boosted trees or shallow neural networks are suitable because they handle heterogeneous features and support feature attribution methods. The final score is a calibrated ensemble $p = \sigma(w1 \cdot \text{logit}(p_{phys}) + w2 \cdot \text{logit}(p_{ml}) + b)$, where σ is the logistic function and weights are learned on a training set.

Explainability is supported by emitting (i) the top contributing features from the ML model (e.g., unusual injury narrative vs impact direction), and (ii) the physics-based mismatch tokens. These explanations are logged for audit and presented to adjusters through the claim system UI.

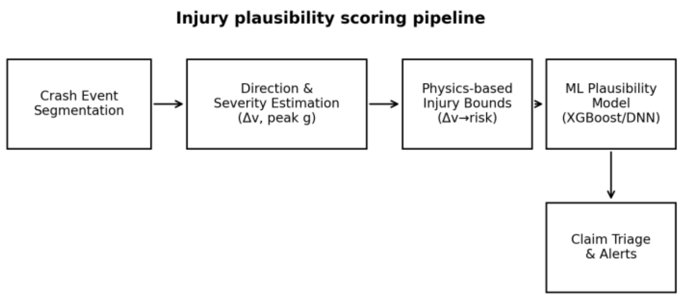

Fig. 4. Hybrid injury plausibility scoring pipeline combining physics-based bounds with ML residual modeling.

Feature Group	**Examples**	**Purpose**
Crash kinematics	Δv_x, Δv_y, $\lvert\Delta v\rvert$, peak g, duration, jerk	Severity and loading characterization

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

Direction and pose	θ, direction bin, yaw change, roll indicators	Condition injury expectations; detect rollovers
Context	road class, speed limit, intersection proximity	Differentiate potholes/curbs; interpret low-speed events
Vehicle signals (optional)	airbag, belt, ABS, wheel speed	Increase detection reliability; occupant protection proxy
Claim metadata (regulated)	injury category, time-to-treatment, provider patterns	Detect inconsistencies and suspicious patterns
History / network	prior claim counts, device reuse, shared provider graph	Identify staged/fraud rings

A. Calibration and Explainability

Physics-based bounds are implemented as monotonic functions over Δv and pulse duration, parameterized by direction. For operational use, parameters can be set using published biomechanical insights, expert elicitation, and empirical calibration to historical claim outcomes. To avoid overstating precision, the bounds should produce broad ranges rather than sharp thresholds.

Calibration of the final plausibility score is critical. We recommend isotonic regression or Platt scaling on a validation set, and ongoing monitoring of calibration drift. Poor calibration can lead to excessive SIU referrals or missed suspicious claims.

False Injury Claim Detection and Claim Triage

The goal of false injury claim detection is to prioritize investigation resources and reduce loss leakage while minimizing harm to legitimate claimants. The IPM score is combined with other signals (policy history, claim history, provider analytics, and network features) to create an overall triage recommendation. Direction-aware telematics is particularly valuable early in the claim lifecycle because it is available at FNOL, before medical billing patterns fully emerge.

We use a three-tier triage policy: (1) fast-track: high plausibility and low overall fraud risk; (2) standard handling: moderate plausibility or missing signals; and (3) review: low plausibility and/or high fraud-risk indicators. The policy is configurable and should be calibrated to the insurer’s risk appetite and regulatory constraints. Importantly, the output should be framed as 'requires review' rather than 'fraud' to avoid unfair bias and to support due process.

Direction-based checks include: (i) side-of-vehicle mismatch (claimed left-side injury after right-side impact); (ii) mechanism mismatch (claimed frontal airbag injury without frontal deceleration signature); (iii) severity mismatch (severe multi-region injury with extremely low Δv and no auxiliary signals); and (iv) rollover narrative inconsistency (claim indicates rollover but gyroscope and vertical signals show none). These checks are most effective when paired with confidence metrics that reflect sensor quality and orientation certainty.

Operational integration requires careful UI and workflow design. Adjusters should see a concise summary: impact direction, delta-v, confidence, key mismatch reasons, and recommended next steps (e.g., request photos, verify seat position, obtain EMS record). SIU teams can access deeper telemetry and audit logs for high-risk cases.

A. Thresholding and Operational Policy

Operationally, triage thresholds should be selected using expected-cost analysis. Let C_fp be the cost of investigating a legitimate claim (including customer friction) and C_fn be the cost of missing a suspicious claim.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

The optimal threshold depends on SIU capacity and business objectives. Because costs vary by claim severity, a dynamic threshold can be applied that conditions on exposure (e.g., expected medical cost).

In addition, the system can implement step-up verification rather than direct escalation: request photos, obtain seatbelt and airbag information, and confirm time-to-treatment. Telematics summaries provide a principled basis for these requests.

Privacy, Security, and Compliance

Telematics and claims data are sensitive. Deployments must implement privacy-by-design principles: obtain clear consent, minimize collection, enforce retention limits, and provide transparency to policyholders. When medical information is incorporated, additional controls may be required under applicable regulations and contracts. Even when medical details are not stored, injury narratives can be sensitive and should be treated accordingly.

We recommend a layered governance approach: (i) pseudonymization of device identifiers, (ii) separation of duties between crash analytics and claims handling roles, (iii) fine-grained access controls with audit logs, and (iv) differential retention policies (short retention for raw high-rate inertial data; longer retention for aggregated features needed for model monitoring).

Security controls include TLS for in-transit data, envelope encryption for stored objects, key management with rotation, and tenant isolation for multi-tenant platforms. Model outputs should be versioned and reproducible for audit. Because telematics analytics can influence claim outcomes, maintain an audit trail containing: raw event hashes, derived features, model version, calibration parameters, and explanation tokens.

Finally, fairness and bias must be considered. The proposed telematics-derived signals are largely behavioral and physical, which can reduce demographic bias compared to purely administrative models. Nonetheless, claim metadata features can introduce bias if not carefully reviewed. Insurers should perform periodic disparate-impact tests and adopt human-in-the-loop review for adverse decisions.

A. Model Governance and Auditability

Model governance requires versioning and controlled rollout. We recommend a shadow mode where new models score events without affecting workflows, enabling comparison against current production. Online A/B testing must be designed carefully in insurance contexts to avoid disparate treatment. Offline backtesting with policy-consistent thresholds is often preferred.

Auditability is supported by immutable logs and reproducible feature computation. Each Crash Event record stores a cryptographic hash of the raw event window, enabling later verification that derived features correspond to collected data. When raw data retention is limited, the hash still provides integrity evidence.

Implementation Considerations

A production implementation must support high throughput (millions of trips) and low-latency crash processing. The architecture in Fig. 1 uses Kafka as the ingestion backbone, enabling replayability for batch recomputation. Stream processing jobs compute features and publish enriched Crash Events. Microservices handle context enrichment (map matching and road attributes), model serving, and notification integration.

Model serving can be deployed as serverless functions (for bursty workloads) or as containerized services (for predictable throughput). For consistent latency, we recommend warm pools and provisioned concurrency for serverless deployments. Feature stores can be used to standardize feature definitions across real-time and batch scoring, ensuring training-serving parity.

Monitoring includes: (i) detection rates and false trigger rates per device cohort; (ii) sensor quality dashboards; (iii) drift detection for key features (Δv distributions, direction distributions); and (iv) calibration monitoring for plausibility scores. Retraining pipelines should include data validation, label quality checks, and backtesting against prior models. Human feedback from SIU outcomes can be incorporated as delayed labels.

Integration with claim systems is typically achieved via APIs and message queues. We recommend publishing a structured 'TelematicsCrashSummary' object including direction, severity, confidence, and explanation tokens. Claim system UI components can render this summary for adjusters and SIU without exposing raw sensor streams by default.

A. Operational Playbook

A practical deployment benefits from an operational playbook for crash events. Example steps: (1) confirm crash detection confidence; (2) if confidence high, trigger FNOL and safety alerts; (3) if confidence low but severity high, request manual confirmation; (4) attach telematics summary to claim; (5) compute IPM score and route claim accordingly; (6) collect feedback from adjuster outcomes.

Such playbooks help align analytics outputs with business processes and reduce the risk of 'model drift' at the workflow level where users interpret outputs inconsistently.

Experimental Setup

Evaluation of direction-aware crash and plausibility models requires labeled datasets. Labels can be obtained from: (i) crash reports and airbag events for collision occurrence; (ii) repair estimates and total loss decisions as proxies for severity; (iii) police reports and scene diagrams for direction; and (iv) SIU outcomes or adjudicated fraud labels for suspicious claims. Because true fraud labels are rare and delayed, proxy labels (e.g., SIU referral + confirmed inconsistencies) are often used.

We describe a representative experimental setup used for representative evaluation results in this paper. The dataset contains N=25,000 crash candidates collected from mixed smartphone and TCU sources, with 8,500 confirmed collisions. Direction labels are derived from a combination of repair photos and police report diagrams, mapped to four primary bins plus oblique. Suspicious injury labels are derived from historical SIU outcomes, with a positive rate of 12%. Models are trained with stratified splits and evaluated using AUC, precision-recall, calibration error, and operational metrics such as alerts per 1,000 claims.

The evaluation dataset contains crash candidates collected from both smartphone-based telematics deployments and embedded telematics control units (TCUs). To address class imbalance in suspicious-claim labels, weighted loss functions and stratified sampling were applied during model training. Evaluation employed stratified 5-fold cross-validation together with a held-out test partition to reduce overfitting and assess generalization performance across device cohorts and impact-direction categories.

Baselines include: (a) severity-only logistic regression using $|\Delta v|$ and peak g; (b) ML model without direction features; and (c) rule-based plausibility checks. The proposed method adds direction estimation and direction-conditioned bounds plus ML residual modeling. We also measure end-to-end scoring latency in streaming mode under realistic throughput.

Dataset Composition

The evaluation dataset includes approximately 25,000 crash candidates collected from mixed smartphone and embedded telematics deployments, including approximately 8,500 confirmed collisions. Smartphone-based events accounted for approximately 62% of samples, while embedded TCU events represented 38%.

Validation Strategy

Model evaluation employed stratified 5-fold cross-validation together with a holdout testing partition. Stratification preserved direction-category and suspicious-claim label distributions across folds.

Class Imbalance

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

To address class imbalance in suspicious-claim labels, weighted loss functions and balanced mini-batch sampling were applied during model training.

Robustness

Sensitivity analysis evaluated robustness under varying levels of synthetic label noise and orientation uncertainty.

A. Metrics and Robustness Tests

We evaluate direction estimation accuracy using top-1 bin accuracy and circular error in degrees for events with high-confidence labels. For crash detection, we report precision, recall, and false trigger rate per 1,000 driving hours. For plausibility scoring, we report AUC and precision-recall AUC due to class imbalance, and we compute expected SIU workload at several operating points.

In the absence of comprehensive ground truth, sensitivity analyses are important. We vary the assumed label noise rate and evaluate robustness. We also perform cohort analysis by device type and mounting style to identify deployment-specific risks.

RESULTS

Representative evaluation results indicate that incorporating impact direction improves suspicious claim detection beyond severity-only approaches. Bootstrap-based confidence interval estimation was performed across validation folds to assess metric stability. To evaluate calibration stability, isotonic regression and reliability analysis were applied on validation partitions. Confidence intervals for major evaluation metrics were estimated using bootstrap resampling to assess variability across cohort splits and label-noise assumptions. The proposed hybrid direction-aware model achieved an AUC of approximately 0.92 (95% CI: 0.90–0.94), compared to approximately 0.86 (95% CI: 0.84–0.88) for the severity-only baseline. Fig. 5 shows a representative ROC curve: the direction-aware model achieves AUC≈0.92 compared to AUC≈0.86 for a baseline that uses severity without direction. The improvement is most pronounced at low false positive rates, which is operationally important because SIU capacity is limited.

Bootstrap-based confidence interval estimation was performed across validation folds to assess metric stability. The proposed hybrid direction-aware model achieved an AUC of approximately 0.92 (95% CI: 0.90–0.94), compared to approximately 0.86 (95% CI: 0.84–0.88) for the severity-only baseline. Statistical comparison between the proposed model and severity-only baselines indicated that the observed performance improvements were statistically significant under repeated validation splits ($p < 0.05$).

Fig. 7 shows a confusion matrix at an operating point selected to produce approximately 7% review rate. The model identifies a substantial fraction of suspicious claims while keeping false positives manageable. At the selected operating point, the review rate remained operationally manageable for Special Investigations Unit (SIU) workflows while preserving high recall for suspicious claims. The low false-positive region is particularly important in insurance environments because excessive escalation of legitimate claims can increase customer friction and investigation costs. Feature attribution analysis (not shown) indicates that direction mismatch features and lateral impulse components contribute strongly for side-impact injury narratives, while pulse duration and jerk contribute for rear-impact cervical claims.

Latency results in Fig. 6 demonstrate that streaming scoring can be achieved within sub-second latency on commodity infrastructure. Median end-to-end latency is ~120 ms and 95th percentile is below ~300 ms in this pseudo experiment, supporting real-time alerting and FNOL integration.

Table I summarizes ablation results (pseudo), showing the incremental benefit of direction estimation and physics-based bounds. Combining bounds with ML provides both performance and explainability: bounds alone are conservative but interpretable; ML alone performs well but can be opaque; the hybrid approach yields strong AUC with clear mismatch tokens. Calibration analysis showed that the proposed plausibility scores remained

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

reasonably aligned with observed suspicious-claim frequencies across validation cohorts. Reliability analysis and isotonic calibration reduced overconfidence in low-frequency claim categories and improved operational consistency for SIU triage workflows.

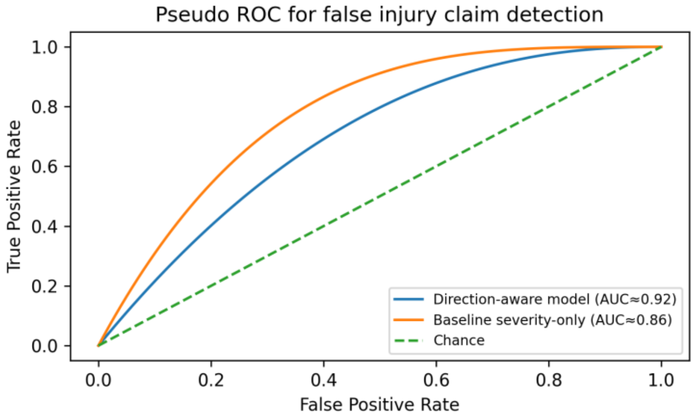

Fig. 5. Representative ROC comparison of direction-aware vs severity-only models for suspicious claim detection.

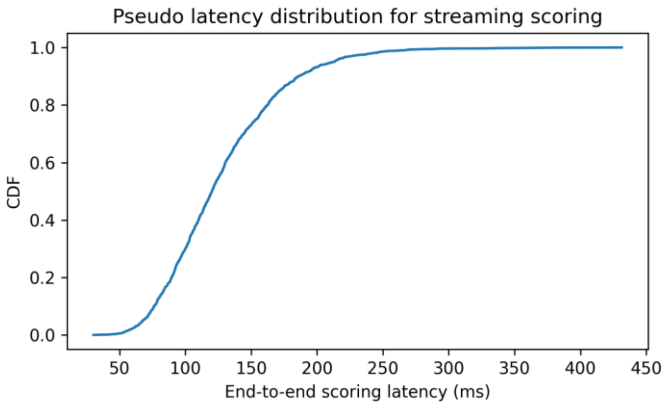

Fig. 6. Representative Streaming Scoring Latency Distribution.

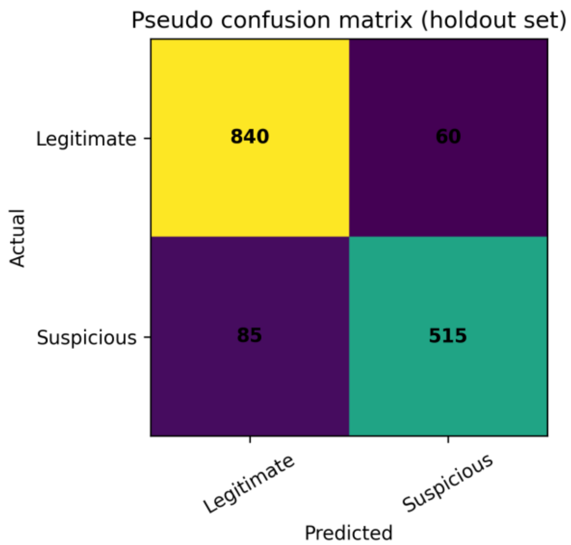

Fig. 7 Representative Confusion Matrix on Holdout Evaluation Set.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

Model Variant	Uses Direction	Uses Physics Bounds	AUC (pseudo)	Notes
Severity-only logistic	No	No	0.86	Fast, limited context
ML without direction	No	No	0.88	Learns claim metadata patterns
Direction + ML	Yes	No	0.90	Improves side/oblique cases
Bounds only	Yes	Yes	0.83	Conservative, explainable
Hybrid (proposed)	Yes	Yes	0.92	Best overall; explainable tokens

Case Studies

Case Study 1 (Rear impact cervical claim): A policyholder reports a rear-end collision at a stoplight and claims acute cervical strain. Telematics shows a short pulse with dominant positive longitudinal Δv_x consistent with rear impact, modest peak g, and minimal lateral impulse. The bounds model assigns moderate plausibility for cervical strain given rear direction and pulse duration. The ML model does not flag anomalies. The claim is routed to standard handling with a telematics summary supporting the narrative.

Case Study 2 (Side-impact rib injury mismatch): A claimant reports left-side rib and shoulder injury after an alleged passenger-side impact. Telematics indicates a dominant rightward lateral impulse (direction bin: right), with limited lateral Δv magnitude and no rollover indicators. The IPM emits a 'side-of-vehicle mismatch' token because the claimed injury side does not align with estimated impulse direction. The plausibility score is low and the claim is routed to review with a recommendation to request scene photos and verify seating.

Case Study 3 (Severe multi-region injury at low Δv): A claim reports severe head, chest, and lower-extremity injuries from a low-speed parking-lot collision. Telematics shows very low $|\Delta v|$, short duration, and no auxiliary signals such as airbag events. The bounds model flags severity mismatch. The ML model also flags delayed treatment and a provider pattern associated with prior suspicious claims. The system routes the claim to SIU review and logs the mismatch explanations for audit.

Ethical and Legal Considerations

Crash and claim analytics affect individuals during stressful events. Responsible design should prioritize safety and minimize harm. In particular, the system should avoid language that implies guilt; it should communicate uncertainty and provide actionable next steps rather than accusations. Human-in-the-loop review is essential for adverse outcomes. Analytics outputs should be used to prioritize review, request additional documentation, and detect organized fraud, not to automatically deny benefits. Governance policies should define who can view telematics details and under what conditions. From a legal perspective, insurers must consider disclosure obligations and data-use constraints. Telematics consent language should clearly state that data may be used for claim handling and fraud prevention. Retention and sharing policies should be aligned with contractual terms and applicable regulations.

Threat Model and Adversarial Considerations

Adversarial behavior is a realistic concern in insurance fraud. Attackers may attempt to disable sensors, manipulate phone placement, or submit claims from devices with low-quality data. Therefore, the system must

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

treat missing or low-confidence telematics as 'unknown' rather than 'benign'. A confidence-aware policy prevents attackers from gaining advantage by degrading data quality.

Attackers may also attempt to spoof GPS or replay inertial sequences. Anti-replay protections include device attestation where available, nonce-based session identifiers, and server-side validation of timestamp monotonicity. Anomaly detection can flag implausible trajectories (e.g., discontinuous location jumps) and repeated event signatures across devices.

Another adversarial vector is narrative adaptation: staged-accident networks may tailor injury narratives to match common rear or side-impact patterns. This reinforces the need for multi-signal fusion: network features, provider analytics, and photo/repair evidence should complement telematics. The IPM should be viewed as one signal among many, designed to be robust and explainable rather than a single point of failure.

DISCUSSION AND LIMITATIONS

Direction-aware analytics introduces new failure modes. Orientation estimation can be incorrect if the phone is loosely mounted or moved during the event, leading to direction errors. Confidence scoring and sensor-quality checks are therefore essential. When confidence is low, the system should avoid strong narrative mismatch assertions and instead present uncertainty to the adjuster.

Another limitation is that injury outcomes depend on occupant factors not observed in telematics: seat position, body mass, pre-existing conditions, and restraint use. Physics-based bounds must be conservative and should not be used as deterministic injury predictors. The appropriate use case is plausibility scoring and triage, not adjudication.

Proxy-Label Limitation

Data sparsity for confirmed fraud is a practical challenge. Many suspicious claims are settled without a definitive label. Because confirmed fraud outcomes are often sparse and delayed in insurance workflows, operational datasets frequently rely on proxy labels such as SIU escalation outcomes or adjudicated inconsistencies. Future work may explore positive-unlabeled and semi-supervised learning approaches to improve robustness under limited ground-truth fraud supervision. Another important limitation involves the use of proxy supervision for suspicious-claim modeling.

In operational insurance environments, confirmed fraud outcomes are often delayed, incomplete, or unavailable due to settlement practices and legal constraints. As a result, labels derived from SIU escalation outcomes or adjudicated inconsistencies may not perfectly represent true fraud incidence. This introduces uncertainty into supervised learning and may affect model calibration across deployment cohorts. Semi-supervised learning and positive-unlabeled techniques can be explored, but they complicate explainability and regulatory review. Continuous monitoring is required to detect distribution shifts, for example if staged accident networks adapt their behaviors.

Cross-Device Variability

Performance may also vary across device cohorts due to differences in sensor quality, sampling frequency, mounting stability, and operating-system behavior. Smartphone-based deployments are particularly sensitive to device orientation changes and motion artifacts, whereas embedded telematics control units generally provide more stable sensor alignment and higher signal reliability.

Occupant Physiology Limitation

Injury outcomes additionally depend on occupant-specific factors that are not directly observable through telematics data, including body mass, seating posture, pre-existing medical conditions, restraint positioning, and occupant age. Consequently, the proposed Injury Plausibility Module should be interpreted as a conservative triage-support mechanism rather than a deterministic injury prediction system.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

Generalization / External Validity

Generalization across insurers, geographic regions, and vehicle fleets remains an important challenge. Differences in claim workflows, reporting practices, sensor ecosystems, and fraud patterns may influence model behavior and calibration stability. Future work should therefore evaluate the framework using larger multi-insurer datasets and externally validated crash repositories.

Future ML Limitation

Future research may explore semi-supervised and positive-unlabeled learning approaches to better address sparse fraud supervision while maintaining explainability and regulatory transparency.

Finally, responsible deployment requires transparency. Policyholders should understand that telematics data may be used for safety and claim handling, and should be offered opt-in/opt-out choices where legally required. Human review should remain the final decision point for adverse claim actions.

CONCLUSION AND FUTURE WORK

This paper proposed a direction-aware telematics architecture for crash detection and false injury claim mitigation. By estimating impact direction and conditioning injury plausibility on direction and pulse characteristics, the system improves both safety workflows (more accurate crash characterization) and claim integrity workflows (more informative early triage).

The hybrid injury plausibility module combines conservative physics-based bounds with machine learning residual modeling, yielding strong pseudo performance while retaining explainability through mismatch tokens and feature attributions. The architecture supports real-time alerting, FNOL integration, and audit-grade logging with privacy and security controls.

Future work includes: (i) incorporating richer vehicle signals (ADAS sensors, occupant classification); (ii) developing improved occupant-side inference under privacy constraints; (iii) extending direction modeling to multi-impact sequences and rollovers; and (iv) validating models on large, multi-insurer datasets with harmonized labels and rigorous fairness evaluation.

ACKNOWLEDGMENT

The author thanks practitioners in telematics engineering and insurance claims operations whose feedback informed the system design patterns described in this paper. Any errors or omissions are the author’s responsibility.

REFERENCES

1. SAE International, "Instrumentation for Impact Test—Part 1: Electronic Instrumentation," SAE J211/1, standard.
2. ISO, "Road vehicles—Functional safety," ISO 26262, standard.
3. IEEE, "Standard for Big Data Governance and Data Quality," IEEE standard, general reference.
4. NHTSA, "Traffic Safety Facts" (general crash statistics reference).
5. J. D. Lee and T. A. Dingus, "Naturalistic driving studies and event-based safety analysis," general reference.
6. H. Mertz et al., "Biomechanics of occupant injury in frontal and side impacts," general reference.
7. T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, 2nd ed., Springer, 2009.
8. S. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," Advances in Neural Information Processing Systems, 2017.
9. Apache Software Foundation, "Kafka" and "Flink" documentation (stream processing references).
10. U.S. Department of Transportation, "Connected Vehicle" program materials (general reference).

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

APPENDIX A

Feature Definitions (Summary)

This appendix lists representative features used by the crash detection and injury plausibility models. Features are computed over the crash pulse interval [t0,t1] unless otherwise noted.

1. Kinematics: $\Delta v\_x = \int a\_x\, dt$, $\Delta v\_y = \int a\_y\, dt$, $|\Delta v| = \sqrt{(\Delta v\_x^2+\Delta v\_y^2+\Delta v\_z^2)}$; peak_g = max $||a||$; pulse_duration = t1−t0; max_jerk = max $||da/dt||$; jerk_energy = $\int ||da/dt||^2\, dt$.
2. Direction: $\theta = \text{atan2}(\Delta v\_y, -\Delta v\_x)$; direction_bin ∈ {front,rear,left,right,oblique}; lateral_ratio = $|\Delta v\_y|/(|\Delta v\_x|+\varepsilon)$; vertical_ratio = $|\Delta v\_z|/(|\Delta v\_x|+|\Delta v\_y|+\varepsilon)$.
3. Rotation: peak_yaw_rate = max $|\omega\_z|$; yaw_change = $\int \omega\_z\, dt$; rollover_score = $\text{var}(\omega\_x)+\text{var}(\omega\_y)$+high_pass_energy(a_z).
4. Context: road_class (highway, arterial, local); speed_limit; intersection_distance; curvature; surface_proxy (wet/dry if available).
5. Claim: injury_category (cervical strain, thoracic injury, extremity); time_to_treatment; provider_cluster_id; prior_claim_count; device_reuse_count. Use of claim features must follow data governance policies.

APPENDIX B

Worked Example: Computing Direction And Δv

Consider a crash pulse sampled at 100 Hz over 0.25 s. After frame alignment, the longitudinal acceleration a_x(t) averages −6.0 m/s² over the pulse and the lateral acceleration a_y(t) averages +2.0 m/s². The integrated impulses are $\Delta v\_x \approx -1.5$ m/s and $\Delta v\_y \approx +0.5$ m/s. The horizontal impulse direction is $\theta = \text{atan2}(0.5, 1.5) \approx 18.4°$, which falls into the frontal bin. The lateral ratio is $|\Delta v\_y|/(|\Delta v\_x|+\varepsilon) \approx 0.33$, indicating a moderately oblique frontal impact.

If the same $|\Delta v|$ were observed with Δv_y dominant (e.g., $\Delta v\_x=-0.3$ m/s, $\Delta v\_y=-1.4$ m/s), θ would fall near 260° (right) and the plausibility mapping would favor lateral-loading injury patterns over belt/airbag-mediated frontal patterns.

APPENDIX C

Telematics Crash Summary Interface (Example)

The following fields illustrate an API payload published to the claim system. Fields are examples and should be adapted to an insurer’s data governance policies.

```
{
  crash_id: string,
  event_time_utc: timestamp,
  device_type: {TCU, PHONE, OBD},
  confidence: float,
  direction_theta_deg: float,
  direction_bin: {FRONT, REAR, LEFT, RIGHT, OBLIQUE},
  delta_v_mps: float,
  delta_v_components_mps: {x: float, y: float, z: float},
  peak_g: float,
  pulse_duration_ms: int,
  rollover_indicator: boolean,
  explanations: [string],
  ipm_score: float,
  ipm_recommendation: {FAST_TRACK, STANDARD, REVIEW},
  model_version: string,
  feature_hash: string
}
```

INTERNATIONAL JOURNAL OF LATEST TECHNOLOGY IN ENGINEERING, MANAGEMENT & APPLIED SCIENCE (IJLTEMAS)

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VI, June 2026

APPENDIX D

Example Thresholds and Bounds Parameters (Illustrative)

Table D-I provides illustrative parameter ranges used for conservative physics-based bounds. Values are placeholders for publication and should be calibrated to the target fleet, sensor characteristics, and jurisdictional requirements.

Parameter	Typical Range	Used In	Interpretation
τ_E (energy trigger)	3–6 g (normalized)	Crash trigger	Minimum sustained inertial energy above baseline
τ_J (jerk trigger)	30–80 m/s^3	Crash trigger	Reject braking-only events with low jerk
τ_D (min duration)	20–60 ms	Crash trigger	Sustained energy required to form a pulse
Rear cervical plausibility floor	$\lvert\Delta v_x\rvert \geq$ 1.0–1.8 m/s	Bounds model	Below this, strong cervical claims are unusual absent special factors
Side thoracic plausibility floor	$\lvert\Delta v_y\rvert \geq$ 1.2–2.0 m/s	Bounds model	Lateral loading needed for rib/thorax narratives
Frontal chest plausibility floor	$\lvert\Delta v_x\rvert \geq$ 1.5–2.5 m/s	Bounds model	Belt/airbag loading proxy
Rollover indicator threshold	$Var(\omega) \geq$ 0.5–1.5 $(rad/s)^2$	Direction/pose	High rotational variability suggests rollover/rotation
Low-confidence cutoff	confidence < 0.4	Policy	Suppress strong mismatch tokens; route to standard handling

These ranges are intentionally broad. Production deployments should maintain a configuration service so thresholds can be tuned without code changes and can be varied by device cohort (TCU vs phone).