Machine Learning-Based Adaptive Restraint Algorithms Using Pre-Crash Sensor Data and Occupant State Estimation: A Comprehensive Review

Ganesh Shete*, Sachin Ratnaparkhi, Rakshit Gummaraju****

DOI: 10.29322/IJSRP.16.02.2026.p17054
https://dx.doi.org/10.29322/IJSRP.16.02.2026.p17054

Paper Received Date: 16th January 2026
Paper Acceptance Date: 18th February 2026
Paper Publication Date: 24th February 2026

Abstract— *This comprehensive review examines machine learning-based adaptive restraint algorithms using pre-crash sensor data and occupant state estimation, analyzing 15 peer-reviewed studies published between 2015 and 2025. The review systematically categorizes research into four thematic areas: wearable airbag systems, simulation-based adaptive restraint design, in-vehicle occupant detection, and system integration frameworks. Key findings include CNN-based wearable systems achieving 95.75% accuracy, radar-based occupant detection attaining 96.88% validation accuracy, and reinforcement learning demonstrating 100× faster convergence in restraint optimization. Critical research gaps are identified in dataset diversity, real-world validation, sensor fusion integration, real-time deployment, explainability, and fairness evaluation. Eight actionable recommendations are provided for advancing the field.*

Index Terms— Machine Learning, Adaptive Restraint Systems, Pre-Crash Sensors, Occupant State Estimation, Convolutional Neural Networks, Automotive Safety, Airbag Deployment, Vehicle Safety Systems

I. INTRODUCTION

Introduction

Automotive safety systems have evolved from passive restraints to intelligent, adaptive protection mechanisms that leverage machine learning (ML) to optimize occupant outcomes. Modern vehicles integrate pre-crash sensors, in-cabin monitoring, and predictive algorithms to tailor restraint deployment based on crash severity, occupant characteristics, and real-time conditions [1], [2]. This paradigm shift addresses the limitations of one-size-fits-all restraint systems, which fail to account for anthropometric diversity, seating posture variations, and dynamic crash scenarios [3], [4].

The integration of ML with pre-crash sensor data enables three critical capabilities: (1) early crash detection and classification using exteroceptive sensors (radar, lidar, cameras), (2) occupant state estimation through in-cabin sensing (radar, cameras, pressure sensors), and (3) adaptive restraint optimization via simulation-driven design and reinforcement learning [5], [6]. Recent advances in deep learning, particularly convolutional neural networks (CNNs), have demonstrated high accuracy in both wearable airbag systems for motorcycles and in-vehicle occupant detection [7], [8]. Simultaneously, surrogate modeling and reinforcement learning have accelerated restraint system optimization, achieving convergence speeds 100× faster than traditional methods [9].

Despite these advances, significant gaps persist in dataset diversity, real-world validation, sensor fusion integration, real-time deployment, and explainability for safety certification [10], [11]. This review systematically analyzes 15 peer-reviewed studies published between 2015 and 2025 to synthesize current methodologies, evaluate performance metrics, identify research gaps, and provide actionable recommendations for advancing adaptive restraint systems toward robust, certifiable deployment.

Research Objectives

This comprehensive review addresses seven specific objectives:

- Systematically categorize ML-based adaptive restraint research into thematic domains (wearable systems, simulation-based design, in-vehicle detection, system integration).
- Analyze ML algorithms employed across studies, including deep learning (CNNs, transfer learning), tree-based methods (XGBoost, Random Forest), surrogate modeling (Gaussian processes), and reinforcement learning.
- Evaluate sensor modalities and data sources, encompassing IMUs, radar, cameras, and simulation platforms (MADYMO, FEM).

- Synthesize occupant state estimation methods, from accident detection in wearable systems to anthropometric classification and posture recognition in vehicles.
- Compare performance metrics across studies, including classification accuracy, convergence speed, and injury risk reduction.
- Identify critical research gaps in dataset diversity, real-world validation, sensor fusion, real-time deployment, and explainability.
- Provide evidence-based recommendations for future research and industry implementation.

Scope and Structure

This review focuses exclusively on ML-based adaptive restraint systems utilizing pre-crash sensor data and occupant state estimation, published between 2015 and 2025. The analysis encompasses wearable airbag systems for motorcycles, simulation-driven restraint design, in-vehicle occupant detection, and integrated system frameworks. Section II presents a thematic literature review with comparative analysis. Section III details methodological approaches across ML techniques, sensor modalities, and evaluation frameworks. Section IV synthesizes results and discusses performance patterns, algorithm effectiveness, and implementation challenges. Section V identifies critical research gaps. Sections VI and VII provide conclusions and actionable recommendations for advancing the field toward certifiable, equitable adaptive restraint systems.

II. LITERATURE REVIEW

Literature Review

This section systematically analyzes 15 peer-reviewed studies, organized into four thematic categories: wearable airbag systems, simulation-based adaptive restraint design, in-vehicle occupant detection, and system integration frameworks. The review synthesizes methodologies, sensor modalities, ML algorithms, and performance metrics to establish the current state of adaptive restraint research.

Wearable Airbag Systems for Motorcycles

Wearable airbag systems represent a critical safety innovation for motorcyclists, who lack the structural protection of enclosed vehicles. These systems employ inertial measurement units (IMUs) and ML algorithms to detect accident situations in real-time and deploy neck-protecting airbags.

Jeong et al. [1] developed a parallel neural network–CNN architecture for a wearable motorcycle airbag system using the MPU6050 six-axis IMU. The system measures acceleration and angular velocity to compute motion angles and classify accident situations. The parallel architecture combines the spatial feature extraction capabilities of CNNs with the temporal processing of standard neural networks, improving judgment performance over single-architecture approaches. Woo et al. [3] demonstrated that CNN-based accident recognition achieves 95.75

These wearable systems share common design principles: (1) reliance on compact, low-power IMUs suitable for wearable integration, (2) supervised learning on labeled accident-type datasets, (3) real-time inference requirements (sub-100ms latency), and (4) emphasis on reducing false positives to prevent unnecessary deployments. However, dataset diversity remains limited, with most studies collecting data from controlled experimental scenarios rather than naturalistic riding conditions.

Simulation-Based Adaptive Restraint Design

Simulation-driven approaches leverage physics-based crash models to generate large-scale datasets for ML-based restraint optimization, enabling exploration of design spaces infeasible for physical testing.

Sun et al. [5] employed Gaussian process surrogate modeling trained on 2,000 MADYMO full-frontal crash simulations to design adaptive restraint policies accounting for occupant anthropometry (sex, stature, BMI). The surrogate model approximates the computationally expensive MADYMO simulator, enabling rapid evaluation of restraint parameter combinations. Population-wise optimization reduced joint injury risk across diverse occupant subgroups, with greatest benefits for tall obese males and short obese females—populations underserved by traditional one-size-fits-all designs. This work demonstrates that adaptive policies can simultaneously improve average protection and reduce disparities across anthropometric subgroups.

Sequeira et al. [6] evaluated six ML classifiers (XGBoost, CatBoost, Random Forest, Kernel SVM, KNN, ANN) for predicting occupant injury criteria from FEM simulation outputs. The task framed injury prediction as binary classification (Safe Zone vs. Injury Zone) to inform pre-crash airbag activation decisions. XGBoost emerged as the best-performing algorithm, leveraging gradient-boosted decision trees to capture nonlinear relationships between crash parameters and injury outcomes. Feature engineering played a critical role, with domain-informed feature selection improving classifier performance.

Mathieu et al. [11] applied reinforcement learning (RL) to optimize five restraint system parameters simultaneously, minimizing occupant loads in frontal impacts. The RL agent achieved approximately 100× faster convergence to comparable optima than global optimization methods, demonstrating the efficiency of policy-gradient approaches for high-dimensional design problems. The trained agent transferred successfully to unseen anthropomorphic test device (ATD) configurations, suggesting learned policies capture generalizable restraint principles rather than overfitting to specific scenarios.

Simulation-based methods enable systematic exploration of design spaces and population-level optimization but face simulation-to-reality gaps. Validation on physical crash tests and integration with online occupant sensing remain critical challenges for deployment.

In-Vehicle Occupant Detection and Classification

In-cabin sensing systems estimate occupant presence, position, and characteristics to enable adaptive restraint deployment and safety feature activation.

The 2022 radar-based study [12] employed FMCW radar and CNN-based transfer learning (VGG-16) for in-vehicle occupant detection. Raw radar time-domain signals were converted to heatmaps, which served as inputs to the CNN. The system achieved 96.88

Perrett et al. [15] developed a cost-based transfer learning framework for occupant detection and classification using a single overhead camera. The system detects and classifies occupants (adult/child) across all passenger seats, enabling automatic child locks and airbag suppression. Cost-sensitive weighting penalizes safety-critical misclassifications (e.g., classifying a child as an adult, leading to inappropriate airbag deployment). Transfer learning leverages data across seats, addressing the challenge of limited labeled data per seating position. The framework demonstrated effectiveness on challenging datasets with varying occupant positions and occlusions.

In-vehicle detection systems must balance accuracy, privacy, computational efficiency, and robustness to environmental variations. Radar-based approaches address privacy concerns but require larger datasets to match camera-based performance. Multi-modal fusion combining radar, cameras, and pressure sensors represents a promising direction for robust occupant state estimation.

System Integration and Predictive Safety Frameworks

Several studies address system-level integration of pre-crash sensing, occupant state estimation, and adaptive restraint control. Zhao et al. [2] proposed an adaptive and proactive safety system integrating sensor data with fuzzy logic for occupant injury prediction. The fuzzy integration approach handles uncertainty in sensor measurements and crash predictions, providing robust decision-making under incomplete information. Sequeira [7] presented a comprehensive framework for prediction-based activation of vehicle safety systems, encompassing crash validation, geometry estimation, and crash severity plus restraint strategy prediction. The work introduced a contact-based validation sensor to complement forward-looking sensors, improving crash detection reliability. Straßburger et al. [10] advocated a top-down systems engineering approach for predictive safety systems, emphasizing the integration of vehicle sensor data with occupant position estimation. This holistic perspective addresses the complexity of coordinating multiple subsystems (perception, prediction, actuation) within stringent real-time and safety constraints.

Kumar et al. [14] proposed optimization and clustering methods to construct lookup tables for adaptive restraint systems, enabling efficient real-time decision-making. ML-assisted clustering structures the design space into discrete regions, each associated with optimal restraint parameters. This approach balances adaptivity with computational efficiency, suitable for embedded deployment on vehicle electronic control units (ECUs).

System integration research highlights the need for robust sensor fusion, fail-safe architectures, and validation frameworks that span perception, prediction, and actuation subsystems. Certification-ready evaluation pipelines and explainability mechanisms remain critical gaps for regulatory approval.

Comparative Analysis

Table tab:comparison provides a comprehensive comparison of the 15 reviewed studies across key dimensions: application domain, ML algorithms, sensor modalities, occupant state estimation methods, and reported performance metrics.

table*[t]

Comprehensive Comparison of Reviewed Studies

tab:comparison

table*

Key patterns emerge from this comparison: (1) Deep learning, particularly CNNs, dominates perception tasks (wearable detection, in-cabin monitoring), (2) Tree-based methods (XGBoost, Random Forest) excel at simulation-based injury prediction, (3) Surrogate modeling and RL enable efficient design optimization, (4) IMUs serve wearable applications while radar and cameras support in-vehicle sensing, and (5) Occupant state estimation ranges from binary accident detection to detailed anthropometric classification. Performance metrics vary widely, reflecting diverse application contexts, but high-performing systems consistently achieve >95

Synthesis of Methodological Trends

Across the reviewed studies, several methodological trends emerge that characterize the evolution of ML-based adaptive restraint research. First, there is a clear progression from single-modality sensing to multi-modal integration, though explicit fusion frameworks remain underexplored. Early wearable systems [1], [3], [4] relied exclusively on IMU data, while recent in-vehicle systems [12], [15] explore radar and camera modalities. However, few studies systematically compare modalities or develop fusion architectures that leverage complementary strengths.

Second, the field exhibits a shift from reactive to predictive safety paradigms. Traditional restraint systems react to crash detection signals, deploying airbags after impact initiation. Pre-crash sensing [7], [10] enables predictive activation based on forward-looking sensors (radar, lidar, camera), allowing restraint systems to prepare for imminent collisions. This temporal shift from milliseconds (post-crash) to hundreds of milliseconds (pre-crash) expands the design space for adaptive protection strategies, including pre-tensioning seatbelts, adjusting headrests, and optimizing airbag inflation profiles.

Third, simulation-driven design has become central to adaptive restraint research, enabling systematic exploration of design spaces infeasible for physical testing. The reviewed studies leverage MADYMO [5] and FEM [6] platforms to generate thousands of crash scenarios, training surrogate models and ML classifiers on synthetic data. This approach accelerates design iteration and enables population-level optimization across anthropometric diversity. However, the simulation-to-reality gap remains a critical challenge, requiring careful validation on physical crash tests and real-world sensor data.

Fourth, there is growing recognition of the importance of population diversity and equity in restraint system design. Sun et al. [5] explicitly optimize for diverse occupant subgroups, demonstrating that adaptive policies can simultaneously improve average protection and reduce disparities. This equity-focused approach addresses historical biases in crash test dummies and restraint designs optimized for average male anthropometry, which underserve women, children, elderly occupants, and individuals with non-average body types.

Finally, the field is transitioning from proof-of-concept demonstrations to deployment-oriented research, though significant gaps remain. Recent studies [11], [12] report convergence speed and validation accuracy metrics relevant to production systems, but comprehensive evaluation of real-time performance, robustness to adversarial inputs, and certification-ready explainability remains limited. Bridging the gap from laboratory prototypes to certifiable, deployable systems requires systematic attention to embedded optimization, sensor fusion, real-world validation, and regulatory compliance.

III. METHODOLOGY

Methodology

This section synthesizes the methodological approaches employed across the reviewed studies, categorizing ML techniques, sensor modalities, evaluation frameworks, and validation strategies. The analysis reveals common patterns and best practices for adaptive restraint system development.

Machine Learning Techniques

The reviewed studies employ five primary categories of ML techniques, each suited to specific tasks within the adaptive restraint pipeline:

Deep Learning (CNNs and Transfer Learning): Convolutional neural networks dominate perception tasks requiring spatial feature extraction from sensor data. Wearable airbag systems [1], [3], [4] employ CNNs to process IMU-derived motion signals, learning discriminative features across accident types. In-vehicle occupant detection [12] applies transfer learning from VGG-16 pre-trained on ImageNet to radar heatmaps, leveraging learned representations to accelerate convergence despite domain shift. Camera-based occupant classification [15] similarly employs transfer learning with cost-sensitive weighting to handle class imbalance and safety-critical misclassifications. CNNs excel at extracting hierarchical features from high-dimensional sensor data but require substantial labeled datasets and computational resources.

Tree-Based Methods (XGBoost, Random Forest, CatBoost): Gradient-boosted decision trees and random forests are preferred for simulation-based injury prediction tasks [6]. These methods handle tabular data with mixed feature types, capture nonlinear interactions, and provide feature importance metrics for interpretability. XGBoost emerged as the best performer for binary injury classification (Safe Zone vs. Injury Zone), leveraging gradient boosting to iteratively refine predictions. Tree-based methods offer computational efficiency, robustness to feature scaling, and inherent handling of missing data, making them suitable for real-time embedded deployment.

Surrogate Modeling (Gaussian Processes): Gaussian process regression [5] approximates computationally expensive physics-based simulators (MADYMO, FEM), enabling rapid evaluation of design alternatives. Trained on 2,000 MADYMO simulations, the surrogate model predicts injury outcomes for arbitrary restraint parameter combinations with quantified uncertainty. This approach enables population-wise optimization across anthropometric subgroups, balancing average performance with equity across diverse occupants. Surrogate modeling is essential for design-space exploration when physical testing or high-fidelity simulation is prohibitively expensive.

Reinforcement Learning: RL agents [11] optimize restraint parameters by interacting with simulation environments, learning policies that minimize occupant loads through trial-and-error. Policy-gradient methods achieve 100× faster convergence than global optimization by exploiting gradient information and learned value functions. RL's ability to optimize multiple parameters simultaneously and transfer to unseen configurations makes it promising for adaptive restraint design, though sample efficiency and safety guarantees during exploration remain challenges.

Fuzzy Logic and Hybrid Approaches: Fuzzy logic integration [2] handles uncertainty in sensor measurements and crash predictions, providing robust decision-making under incomplete information. Hybrid approaches combining ML with rule-based logic enable interpretable, certifiable systems that leverage both data-driven learning and domain expertise.

Sensor Modalities and Data Sources

Adaptive restraint systems integrate diverse sensor modalities, each with distinct advantages and limitations. **Inertial Measurement Units (IMUs)** provide high-frequency acceleration and angular velocity measurements suitable for wearable applications [1], [3], [4]. The MPU6050 six-axis IMU is widely adopted for motorcycle airbag systems due to its compact form factor, low power consumption, and sufficient sampling rate (typically 100-200 Hz) for accident detection. **FMCW Radar** enables privacy-preserving in-cabin occupant detection [12], generating heatmaps from time-domain signals that encode occupant presence and position. Radar operates reliably across lighting conditions and does not capture identifiable images, addressing privacy concerns. **Cameras** (overhead, dashboard-mounted) support detailed occupant classification [15], including adult/child discrimination and posture estimation, but raise privacy concerns and require robust performance across lighting variations. **Simulation Platforms** (MADYMO, FEM) generate large-scale synthetic datasets [5], [6], [11] for training and optimization, enabling systematic exploration of crash scenarios and occupant anthropometry infeasible for physical testing. However, simulation-to-reality gaps necessitate careful validation on physical crash tests.

Evaluation Frameworks

Performance evaluation varies across application contexts. **Classification Accuracy** is the primary metric for occupant detection and accident recognition tasks [3], [12], [15], with top-performing systems achieving 95-97

Validation Approaches

Validation strategies span simulation-based testing, controlled experiments, and limited field deployment. Most studies rely on **simulation validation** [5], [6], [11], comparing ML predictions against ground-truth simulation outputs. **Experimental validation** [1], [3], [4] collects data from controlled accident scenarios (e.g., motorcycle drops, staged collisions) to train and test wearable systems. **Dataset-based validation** [12], [15] evaluates models on held-out test sets, reporting accuracy, precision, recall, and confusion matrices. However, **real-world field validation** remains limited, with few studies reporting performance on naturalistic driving or riding data. The simulation-to-reality gap, dataset diversity, and long-tail scenario coverage represent critical validation challenges for certification and deployment.

Data Collection and Annotation Strategies

The reviewed studies employ diverse data collection methodologies, each with distinct trade-offs between cost, realism, and scalability. Wearable airbag research [1], [3], [4] collects IMU data through controlled experimental protocols, including motorcycle drops from various heights and angles, staged low-speed collisions, and instrumented test rides. These controlled scenarios enable precise labeling of accident types and ground-truth timing for deployment decisions, but may not capture the full diversity of real-world crash dynamics, rider behaviors, and environmental conditions (road surface, weather, traffic).

In-vehicle occupant detection studies [12], [15] collect sensor data (radar heatmaps, camera images) in stationary or low-speed driving conditions, systematically varying occupant positions, seating configurations, and vehicle types. This approach enables efficient data collection with precise ground-truth labels (occupant presence, seat location, adult/child classification) but may not reflect dynamic driving conditions, occupant motion during maneuvers, or sensor performance degradation over vehicle lifetime.

Simulation-based studies [5], [6], [11] leverage design-of-experiments (DoE) methodologies to systematically sample parameter spaces, including crash severity (impact speed, angle), occupant anthropometry (sex, stature, BMI, age), seating position (seat track, seatback angle), and restraint parameters (airbag inflation timing, seatbelt pretensioner force). Latin hypercube sampling, maximal projection designs, and adaptive sampling strategies enable efficient coverage of high-dimensional design spaces with limited simulation budgets (typically hundreds to thousands of runs). However, simulation fidelity depends on validated material models, contact algorithms, and human body models, requiring careful verification against physical crash tests.

Annotation strategies vary from automated labeling (simulation ground truth, sensor-derived metrics) to manual expert annotation (occupant posture classification, crash severity assessment). Inter-annotator agreement, label noise, and class imbalance represent common challenges, particularly for rare events (severe crashes, out-of-position occupants) and ambiguous cases (borderline injury thresholds, occluded occupants). Active learning and semi-supervised learning techniques offer promising approaches to reduce annotation burden while maintaining label quality, but remain underexplored in the reviewed literature.

Model Training and Hyperparameter Optimization

Training procedures across studies reflect diverse objectives and computational constraints. Wearable airbag systems [1], [3], [4] employ supervised learning with cross-entropy loss, training CNNs for 50-200 epochs using Adam or SGD optimizers with learning rate schedules. Data augmentation (rotation, scaling, noise injection) improves generalization to unseen accident scenarios. Early stopping based on validation loss prevents overfitting to limited training data.

In-vehicle detection systems [12], [15] leverage transfer learning, fine-tuning pre-trained models (VGG-16, ResNet) on domain-specific data. Transfer learning accelerates convergence (10-50 epochs vs. 100+ epochs for training from scratch) and improves generalization when labeled data is limited. Layer-wise learning rates (lower rates for early layers, higher rates for task-specific layers) preserve learned representations while adapting to new domains.

Simulation-based surrogate models [5] employ Gaussian process regression with carefully selected kernel functions (squared exponential, Matérn) and hyperparameters (length scales, noise variance) optimized via maximum likelihood estimation. Surrogate model accuracy is validated through leave-one-out cross-validation and comparison against held-out simulation runs, ensuring reliable approximation of expensive simulators.

Reinforcement learning agents [11] train through interaction with simulation environments, using policy gradient methods (PPO, TRPO) or value-based methods (DQN, SAC). Reward shaping (combining injury metrics, deployment timing, and constraint violations) guides learning toward safe, effective policies. Curriculum learning (progressively increasing crash severity, occupant diversity) improves sample efficiency and convergence stability. Hyperparameter optimization (learning rate, discount factor, network architecture) employs grid search, random search, or Bayesian optimization, though computational cost often limits systematic tuning.

Despite diverse training methodologies, few studies report comprehensive ablation studies isolating the impact of architectural choices, data augmentation strategies, or hyperparameter settings. Systematic reporting of training procedures, computational requirements (GPU-hours, wall-clock time), and sensitivity to hyperparameters would improve reproducibility and enable fair comparison across methods.

IV. RESULTS AND DISCUSSION

Results and Discussion

This section synthesizes key findings across the reviewed studies, identifying common performance patterns, algorithm effectiveness, and implementation challenges. The analysis reveals both the promise of ML-based adaptive restraint systems and the barriers to widespread deployment.

Common Performance Patterns

Three consistent patterns emerge across high-performing systems. First, **deep learning architectures achieve superior accuracy on perception tasks** when sufficient labeled data is available. CNN-based wearable airbag systems [3] attain 95.75

Second, **simulation-driven optimization enables population-level adaptation** that traditional one-size-fits-all designs cannot achieve. Gaussian process surrogate modeling [5] trained on 2,000 MADYMO simulations enables adaptive restraint policies that reduce injury risk across diverse anthropometric subgroups, with greatest benefits for underserved populations (tall obese males, short obese females). This population-wise approach addresses equity concerns, ensuring that safety systems protect all occupants rather than optimizing for average anthropometry. The surrogate modeling framework balances computational efficiency with design-space exploration, enabling rapid evaluation of thousands of parameter combinations.

Third, **reinforcement learning accelerates design optimization** by orders of magnitude compared to traditional methods. RL agents [11] achieve approximately 100× faster convergence to comparable optima than global optimization, leveraging learned value functions and policy gradients to efficiently navigate high-dimensional design spaces. The trained policies transfer successfully to unseen ATD configurations, suggesting that RL captures generalizable restraint principles. This efficiency gain is critical for iterative design processes and real-time adaptive control, though sample efficiency during training and safety guarantees during exploration remain active research challenges.

Algorithm Effectiveness by Task

Table tab:performance summarizes reported performance metrics across studies, organized by application task.

Perception tasks (accident detection, occupancy classification) favor deep learning when labeled data is abundant. CNNs learn hierarchical representations that generalize across sensor modalities (IMU, radar, camera), achieving >95

Prediction and decision tasks (injury classification, restraint activation) benefit from tree-based methods [6] that handle tabular simulation data, provide feature importance for interpretability, and offer computational efficiency for real-time inference. XGBoost's gradient boosting iteratively refines predictions, capturing complex nonlinear relationships between crash parameters and injury outcomes. Tree-based methods are well-suited for embedded deployment on vehicle ECUs with limited computational resources.

Design optimization tasks leverage surrogate modeling [5] and reinforcement learning [11] to efficiently explore high-dimensional design spaces. Surrogate models approximate expensive simulators, enabling rapid evaluation and population-wise optimization. RL agents learn adaptive policies through interaction, achieving faster convergence than gradient-free global optimization. Both approaches require careful validation to ensure learned models and policies generalize beyond training distributions.

Implementation Challenges

Despite promising research results, several implementation challenges hinder widespread deployment of ML-based adaptive restraint systems.

Dataset Diversity and Real-World Validation: Most studies rely on limited experimental datasets or simulation data, with insufficient reporting of performance across diverse real-world conditions. Wearable airbag systems [1], [3], [4] collect data from controlled accident scenarios rather than naturalistic riding, raising concerns about generalization to real-world crash dynamics, environmental conditions, and rider behaviors. In-vehicle detection systems [12] train on datasets with 1,000 samples per class—modest by deep learning standards—and lack evaluation on diverse vehicle types, seating configurations, and occupant demographics. Simulation-based methods [5], [6], [11] face simulation-to-reality gaps, where models trained on MADYMO or FEM outputs may not transfer reliably to physical crash tests due to modeling assumptions, parameter uncertainties, and unmodeled dynamics.

Sensor Fusion and Multi-Modal Integration: While individual sensor modalities demonstrate strong performance, explicit multi-sensor fusion frameworks integrating exteroceptive (radar, lidar, camera) and in-cabin (radar, camera, pressure) sensing are underexplored. Robust occupant state estimation requires fusing complementary modalities: radar for privacy-preserving presence detection, cameras for detailed posture estimation, and pressure sensors for contact-based validation. Fusion architectures must handle sensor failures, conflicting measurements, and varying latencies while maintaining real-time performance and fail-safe operation.

Real-Time Deployment and Computational Constraints: Few studies report latency, memory footprint, or energy consumption metrics critical for embedded deployment on vehicle ECUs. Wearable systems require sub-100ms inference latency and milliwatt-scale power consumption to enable battery operation. In-vehicle systems must process sensor streams at 10-30 Hz while sharing computational resources with other safety-critical functions (ABS, ESC, ADAS). Model compression techniques (quantization, pruning, knowledge distillation) and hardware acceleration (GPUs, TPUs, FPGAs) are essential for meeting real-time constraints but remain underreported in the reviewed literature.

Explainability and Safety Certification: Black-box ML models (deep networks, RL agents) lack the interpretability required for safety certification and regulatory approval. Certification standards (ISO 26262, SOTIF) demand transparent, verifiable decision-making with quantified failure modes and uncertainty estimates. Explainable AI techniques (attention mechanisms, saliency maps, SHAP values) can provide post-hoc interpretability, but formal verification methods for neural networks remain computationally intractable for production-scale models. Hybrid architectures combining ML perception with rule-based decision logic offer a path toward certifiable systems, but require careful co-design to preserve both performance and interpretability.

Cross-Domain Insights and Transferability

Analysis across the four thematic categories reveals opportunities for cross-domain knowledge transfer. Wearable airbag systems [1], [3], [4] demonstrate that CNN architectures can achieve high accuracy on IMU-derived motion signals with relatively small datasets (hundreds to thousands of samples). This finding suggests that transfer learning from wearable systems to in-vehicle IMU-based crash detection could accelerate development of vehicle-integrated pre-crash sensing. Similarly, the radar-based occupant detection work [12] demonstrates that transfer learning from ImageNet-pretrained models (VGG-16) to radar heatmaps achieves strong performance despite significant domain shift. This result indicates that large-scale pre-training on natural images provides generalizable feature representations applicable to automotive sensing modalities.

Simulation-driven optimization techniques [5], [11] developed for in-vehicle restraint design could be adapted to wearable airbag systems, enabling population-level optimization of deployment thresholds and inflation profiles across diverse rider anthropometry. Conversely, the real-time embedded deployment constraints addressed in wearable systems (sub-100ms latency, milliwatt-scale power) provide valuable lessons for in-vehicle systems targeting pre-crash activation with stringent timing requirements.

The cost-sensitive transfer learning framework [15] developed for multi-seat occupant classification addresses a challenge common across safety-critical ML applications: handling class imbalance and asymmetric misclassification costs. This methodology is directly applicable to wearable accident detection (where false negatives are more costly than false positives) and pre-crash activation (where unnecessary deployments cause user distrust but missed activations risk injury). Systematic application of cost-sensitive learning across adaptive restraint domains could improve safety-critical decision-making while maintaining acceptable false positive rates.

Finally, the systems engineering perspectives [7], [10] emphasize the importance of holistic integration across perception, prediction, and actuation subsystems. These frameworks provide architectural guidance applicable to both wearable and in-vehicle systems, highlighting the need for fail-safe operation, sensor redundancy, and graceful degradation under component failures. Adopting systems-level design methodologies early in development can prevent integration challenges that emerge when optimizing individual subsystems in isolation.

V. RESEARCH GAPS

Research Gaps and Future Directions

This review identifies six critical research gaps that must be addressed to advance ML-based adaptive restraint systems from laboratory prototypes to certifiable, deployable safety technologies.

Gap 1: Dataset Diversity and Scale. Current studies rely on limited experimental datasets (hundreds to thousands of samples) or simulation data with insufficient diversity across crash scenarios, occupant demographics, environmental conditions, and vehicle types [3], [4], [12]. Real-world deployment requires datasets spanning: (1) diverse crash modes (frontal, side, rear, rollover, oblique), (2) occupant anthropometry (sex, age, stature, BMI, pregnancy), (3) seating postures (nominal, out-of-position, reclined), (4) environmental conditions (lighting, temperature, vibration), and (5) vehicle configurations (sedan, SUV, truck, electric). Publicly available benchmark datasets with standardized evaluation protocols are needed to enable reproducible research and fair algorithm comparison.

Gap 2: Simulation-to-Reality Transfer. Simulation-based methods [5], [6], [11] demonstrate strong performance on MADYMO and FEM outputs but lack validation on physical crash tests. Bridging the simulation-to-reality gap requires: (1) domain adaptation techniques to align simulation and real-world sensor distributions, (2) uncertainty quantification to identify when models extrapolate beyond training data, (3) hybrid simulation-physical validation pipelines, and (4) transfer learning from simulation to limited real-world data. Systematic reporting of sim-to-real transfer performance is essential for assessing deployment readiness.

Gap 3: Multi-Modal Sensor Fusion. While individual modalities (IMU, radar, camera) achieve high accuracy, explicit fusion frameworks integrating complementary sensors are underexplored [1], [5], [12]. Robust occupant state estimation requires fusing: (1) exteroceptive sensors (radar, lidar, camera) for pre-crash detection, (2) in-cabin sensors (radar, camera, pressure) for occupant classification, and (3) vehicle dynamics (IMU, wheel speed, steering angle) for crash severity estimation. Fusion architectures must handle sensor failures, asynchronous measurements, and conflicting information while maintaining real-time performance and fail-safe operation.

Gap 4: Real-Time Embedded Deployment. Few studies report latency, memory, or energy metrics for embedded deployment on vehicle ECUs [11], [12]. Production systems require: (1) sub-10ms inference latency for pre-crash activation, (2) <100MB memory footprint for resource-constrained ECUs, (3) milliwatt-scale power for wearable systems, and (4) deterministic worst-case execution time for safety-critical scheduling. Model compression (quantization, pruning), hardware acceleration (GPUs, TPUs), and efficient neural architectures (MobileNet, EfficientNet) must be systematically evaluated for automotive deployment.

Gap 5: Explainability and Certification. Black-box ML models lack the interpretability required for ISO 26262 and SOTIF certification [5], [11], [15]. Certifiable systems require: (1) transparent decision-making with human-understandable rationales, (2) quantified uncertainty estimates for out-of-distribution inputs, (3) formal verification of safety properties (e.g., no false negatives for severe crashes), and (4) failure mode and effects analysis (FMEA) for ML components. Hybrid architectures combining ML perception with rule-based decision logic, along with explainable AI techniques (attention, SHAP), offer paths toward certification but require validation on production-scale systems.

Gap 6: Equity and Subgroup Performance. While Sun et al. [5] explicitly address population diversity, most studies lack systematic evaluation across occupant subgroups (sex, age, BMI, ethnicity). Equitable safety systems require: (1) balanced datasets representing diverse populations, (2) fairness metrics quantifying performance disparities, (3) adaptive policies that improve protection for underserved groups, and (4) regulatory frameworks mandating subgroup reporting. Addressing historical biases in crash test dummies and simulation models is essential for equitable adaptive restraint systems.

VI. CONCLUSION

Conclusion

This comprehensive review systematically analyzed 15 peer-reviewed studies on ML-based adaptive restraint systems published between 2015 and 2025, synthesizing methodologies, performance metrics, and research gaps across wearable airbag systems, simulation-driven design, in-vehicle occupant detection, and system integration frameworks. The analysis reveals significant progress in applying deep learning, tree-based methods, surrogate modeling, and reinforcement learning to optimize occupant protection, with top-performing systems achieving >95

Key findings demonstrate that CNN-based architectures excel at extracting discriminative features from high-dimensional sensor data (IMU, radar, camera), enabling accurate accident detection and occupant classification. Simulation-driven approaches leverage Gaussian process surrogates and RL agents to efficiently explore design spaces and optimize adaptive restraint policies across diverse occupant populations, addressing equity concerns inherent in one-size-fits-all systems. Tree-based methods (XGBoost, Random Forest) provide interpretable, computationally efficient injury prediction suitable for real-time embedded deployment.

However, critical gaps persist in dataset diversity, real-world validation, sensor fusion, real-time deployment, explainability, and equity evaluation. Most studies rely on limited experimental or simulation data, with insufficient reporting of performance across diverse crash scenarios, occupant demographics, and environmental conditions. The simulation-to-reality gap remains a fundamental challenge for methods trained on MADYMO or FEM outputs. Multi-modal sensor fusion frameworks integrating exteroceptive and in-cabin sensing are underexplored, despite their potential for robust occupant state estimation. Real-time deployment metrics (latency, memory, energy) are rarely reported, hindering assessment of production readiness. Black-box ML models lack the interpretability required for safety certification under ISO 26262 and SOTIF standards.

Advancing ML-based adaptive restraint systems from research prototypes to certifiable, deployable technologies requires systematic efforts to address these gaps through diverse benchmark datasets, sim-to-real transfer techniques, multi-modal fusion architectures, embedded optimization, explainable AI methods, and equity-focused evaluation frameworks. The recommendations in Section VII provide actionable guidance for researchers and practitioners to accelerate progress toward robust, certifiable, and equitable adaptive restraint systems that protect all occupants across diverse crash scenarios.

VII. RECOMMENDATIONS

Recommendations

Based on the systematic analysis of research gaps and implementation challenges, this review provides eight actionable recommendations for advancing ML-based adaptive restraint systems toward robust, certifiable deployment:

- **Establish Public Benchmark Datasets:** Develop and release diverse, standardized datasets spanning crash modes, occupant demographics, seating postures, environmental conditions, and vehicle types. Include sensor data (IMU, radar, camera), ground-truth labels, and evaluation protocols to enable reproducible research and fair algorithm comparison.
- **Develop Sim-to-Real Transfer Frameworks:** Systematically validate simulation-trained models on physical crash tests using domain adaptation, uncertainty quantification, and hybrid validation pipelines. Report transfer performance metrics to assess deployment readiness and identify failure modes.
- **Integrate Multi-Modal Sensor Fusion:** Design fusion architectures combining exteroceptive (radar, lidar, camera) and in-cabin (radar, camera, pressure) sensors with fail-safe operation, sensor fault detection, and real-time performance. Leverage complementary modalities to improve robustness and reduce false positives/negatives.
- **Optimize for Embedded Deployment:** Apply model compression (quantization, pruning, distillation), hardware acceleration (GPUs, TPUs), and efficient architectures (MobileNet, EfficientNet) to meet real-time constraints (<10ms latency, <100MB memory). Report deployment metrics systematically.
- **Advance Explainable and Certifiable AI:** Develop hybrid architectures combining ML perception with rule-based decision logic, integrate explainable AI techniques (attention, SHAP), and pursue formal verification methods for safety-critical components. Align with ISO 26262 and SOTIF certification requirements.
- **Evaluate Equity and Subgroup Performance:** Systematically report performance across occupant subgroups (sex, age, BMI, ethnicity) using fairness metrics. Design adaptive policies that improve protection for underserved populations and address historical biases in crash test dummies and simulation models.
- **Conduct Large-Scale Field Validation:** Deploy instrumented vehicles and wearable systems to collect naturalistic driving/riding data across diverse conditions. Validate ML models on real-world data to assess generalization, identify long-tail failure modes, and refine algorithms iteratively.
- **Foster Industry-Academia Collaboration:** Establish partnerships to share proprietary datasets, validation facilities, and domain expertise. Develop pre-competitive research consortia to accelerate progress on common challenges (datasets, benchmarks, certification frameworks) while enabling competitive differentiation on algorithms and system integration.

These recommendations provide a roadmap for researchers, practitioners, and policymakers to advance ML-based adaptive restraint systems from laboratory prototypes to production-ready technologies that enhance occupant safety across diverse populations and crash scenarios.

REFERENCES

References
thebibliography99
jeong2020
J.-H. Jeong, S.-H. Jo, J. Woo, D.-H. Lee, and T.-K. Sung, ``Parallel neural network–convolutional neural networks for wearable motorcycle airbag system," *Journal of Electrical*

Engineering & Technology, vol. 15, 2020. DOI: https://doi.org/10.1007/S42835-020-00507-510.1007/S42835-020-00507-5
zhao
B. Zhao, M. Yang, D. Zeng, Y. Cai, and G. Guo, ``Research on adaptive and proactive safety system based on occupant injury prediction."
woo2020
J. Woo, S.-H. Jo, J.-H. Jeong, M. Kim, and G. S. Byun, ``A study on wearable airbag system applied with convolutional neural networks for safety of motorcycle," *Journal of Electrical Engineering & Technology*, vol. 15, 2020. DOI: https://doi.org/10.1007/S42835-020-00353-510.1007/S42835-020-00353-5
jo2019
S.-H. Jo, J. Woo, J.-H. Jeong, and G. S. Byun, ``Safety air bag system for motorcycle using parallel neural networks," *Journal of Electrical Engineering & Technology*, vol. 14, 2019. DOI: https://doi.org/10.1007/S42835-019-00229-310.1007/S42835-019-00229-3
sun2023
W. Sun, J. Liu, J. Hu, J. Jin, and K. Siasoco, ``Adaptive restraint design for a diverse population through machine learning," *Frontiers in Public Health*, vol. 11, 2023. DOI: https://doi.org/10.3389/fpubh.2023.120297010.3389/fpubh.2023.1202970
sequeira2021
G. J. Sequeira, R. Lugner, T. Brandrneier, E. Elnagdy, and G. Danapal, ``Investigation of different classification algorithms for predicting occupant injury criterion to decide the required restraint strategy," in *Proc. IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2021. DOI: https://doi.org/10.1109/ITSC48978.2021.956486810.1109/ITSC48978.2021.9564868
sequeira2022
G. J. Sequeira, ``Prediction based activation of vehicle safety systems - a contribution to improve to occupant safety by validation of pre-crash information and crash severity plus restraint strategy prediction," 2022. DOI: https://doi.org/10.51202/978318681712910.51202/9783186817129
schneider
B. Schneider, D. Kofler, and G. A. D'Addetta, ``Approach for machine learning based design of experiments for."
verma
A. K. Verma, S. Majee, R. Yadav, and S. Mandal, ``Advancements in airbag deployment technology for enhanced passenger safety."
strassburger
P. Straßburger, B. Grotz, S. Heiß, and B. Arbter, ``Predictive safety: towards holistic top-down systems engineering for pre-crash systems."
mathieu2024
J. Mathieu, P. Gupta, M. Di Roberto, and M. Vielhaber, ``Minimizing occupant loads in vehicle crashes through reinforcement learning-based restraint system design: assessing performance and transferability," *Proceedings of the Design Society*, vol. 4, 2024. DOI: https://doi.org/10.1017/pds.2024.21610.1017/pds.2024.216
radar2022
``Novel radar based in-vehicle occupant detection using convolutional neural networks," in *Proc. IEEE Symposium Series on Computational Intelligence (SSCI)*, 2022. DOI: https://doi.org/10.1109/ssci51031.2022.1002214610.1109/ssci51031.2022.10022146
laakmann
F. Laakmann and R. Schöneburg, ``Occupant and partner protection."
kumar
R. D. Ashok Kumar and P. Anand, ``An optimization and cluster based approach to lookup tables in design of adaptive restraint systems."
perrett2015
T. Perrett, M. Mirmehdi, and E. Dias, ``Cost-based feature transfer for vehicle occupant classification," *arXiv: Computer Vision and Pattern Recognition*, 2015.
thebibliography