Deep Learning-Based Risk Classification for Gap Acceptance Decisions at Unsignalized Intersections Toward Proactive V2X Safety System

Insan Arafat Jahan
Department of Civil, Environmental and Construction Engineering, University of Central Florida
Orlando, Florida, 32816
Email: in448736@ucf.edu
ORCID: 0009-0003-5254-6021

Word Count: 3746 words + 2 table(s) × 250 = 4246 words

Submission Date: May 8, 2026

This work has not been peer-reviewed. More experiments are being conducted toward findings substantial enough to yield a journal submission.

ABSTRACT

Unsignalized intersections account for a disproportionate share of traffic fatalities, with gap acceptance misjudgments as a primary crash mechanism. Existing collision warning systems react after conflict has materialized, relying on threshold-based metrics such as time-to-collision that offer little predictive capacity before a crossing is committed. This study presents a machine learning pipeline to classify accepted gap acceptance events as safe or safety-critical using pre-decision trajectory data alone, supporting proactive V2X warning system design. Crossing events were extracted from two geometrically distinct scenarios in the INTERACTION naturalistic driving dataset – a T-junction (EP0) and an all-way stop (MA) – and labeled using Post-Encroachment Time (PET $\leq$ 1.0 s) as the risk criterion, yielding 1,023 events. A 30-frame pre-entry observation window capturing ego kinematics, conflict vehicle kinematics, and pairwise interaction features served as model input. Five classifiers were trained on MA and evaluated on EP0 to assess cross-scenario generalization. The LSTM achieved the best overall performance (AUC = 0.647, SC recall = 0.862), with conflict vehicle geometry identified as the dominant risk predictor and temporal importance rising toward the entry point. A pronounced domain shift between scenarios established that intersection geometry critically shapes the risk signal, motivating scenario-aware training strategies for real-world deployment.

Keywords: Traffic Safety, Surrogate Safety, Gap Acceptance, Post-Encroachment Time, Deep Learning

INTRODUCTION

Intersections account for roughly one-quarter of all traffic fatalities and approximately half of all traffic injuries in the United States (*1*), with nearly two-thirds of intersection fatalities occurring at unsignalized, stop-controlled locations. These are not edge-case scenarios – unsignalized intersections are the most common intersection type in the country, embedded in the everyday fabric of urban and suburban road networks. Yet they operate without the structured conflict resolution that signals provide, placing the full burden of safe crossing decisions on the driver.

At the core of these crashes is the gap acceptance decision – a driver's choice to cross or merge into a conflicting traffic stream. The decision unfolds in seconds, under uncertainty, and with consequences that are largely irreversible. A driver who misjudges the available gap commits to a crossing that the geometry of the situation cannot accommodate. By the time a conflict is detectable by conventional sensing, the margin for intervention has already collapsed.

Current collision warning systems are ill-suited to this problem. They respond to conflict after it has materialized, relying on threshold-based metrics such as time-to-collision (TTC) computed in real time as vehicles converge. This reactive posture limits intervention to the final moments before impact and generates frequent false alarms on events that resolve safely – undermining driver trust and system utility. What is needed is not a faster reaction, but an earlier one: a system capable of assessing whether a gap acceptance decision is likely to produce a dangerous outcome *before* the vehicle commits to the crossing.

With the emergence of Vehicle-to-Everything (V2X) communication and smart intersection infrastructure, this shift is no longer hypothetical. Roadside sensors and connected vehicle data make it feasible to observe pre-crossing behavior at scale, creating the conditions for proactive, trajectory-informed risk classification. This study develops a machine learning pipeline to realize that capability – classifying accepted gap acceptance events as safe or safety-critical using pre-decision trajectory data extracted from the INTERACTION naturalistic driving dataset. Post-Encroachment Time (PET) serves as the ground-truth risk criterion, and five classifiers – ranging from logistic regression to LSTM – are trained on one intersection scenario and evaluated on another to assess cross-scenario generalizability. The central finding is that intersection geometry critically shapes the risk signal, and that models trained on one intersection type do not readily transfer to another – a result with direct implications for how data collection and model deployment strategies must be structured in real-world V2X safety systems.

RELATED WORK

Gap acceptance behavior at unsignalized intersections has been studied extensively through critical gap estimation and behavioral modeling. Li et al. (*2*) used SHRP2 Naturalistic Driving Study data across 64 T-intersections to identify gap size, waiting time, and major-road traffic volume as the dominant predictors of crossing decisions, reporting critical gaps of 5.25–6.19 seconds. Nagalla et al. (*3*) demonstrated that Support Vector Machines, Random Forests, and Decision Trees can effectively classify accept/reject decisions from video-extracted kinematic features. Zgonnikov et al. (*4*) further showed that these decisions are governed by noisy evidence accumulation, with gap acceptance probability driven by time-to-arrival rather than distance — a finding that informed our feature design. Critically, most of this work focuses on predicting whether a crossing maneuver is accepted, rather than evaluating the safety consequence of the accepted crossing. Schumann et al. (*5*) benchmark state-of-the-art gap acceptance prediction models and find that all evaluated models fail precisely in the most safety-critical situations, motivating a shift from predicting decisions to

classifying their risk outcomes.

The theoretical foundation for risk labeling comes from surrogate safety analysis. Hydén (*6*) introduced Post-Encroachment Time (PET) as a continuous, observable proxy for collision proximity, and Laureshyn et al. (*7*) formalized a severity hierarchy for traffic encounters derived from micro-level trajectory data. PET has also been used in more recent trajectory-based crash-risk frameworks, including Bayesian extreme-value approaches for pedestrian safety under heterogeneous traffic conditions by Anowar et al. (*8*). The present work operationalizes these concepts by using PET $\leq$ 1.0 s as the safety-critical threshold. Prior work by the authors has examined how sensing infrastructure affects surrogate safety measurement quality at intersections (*9*), alongside broader contributions to infrastructure-based vehicle detection and road safety monitoring systems (*10–12*), collectively motivating the need for trajectory-level risk classification at the decision point itself.

High-quality naturalistic trajectory data is a prerequisite for this kind of analysis. Krajewski et al. (*13*) established drone-based aerial recording as the standard for naturalistic vehicle trajectory collection, producing over 110,000 vehicle trajectories from German highways. The INTERACTION dataset (*14*) extended this to interactive urban scenarios across the USA, Germany, and China, explicitly including adversarial and near-collision behavior — making it well-suited for safety-focused analysis. For multi-agent interaction modeling, Diehl et al. (*15*) showed that representing a traffic scene as a graph of interacting vehicles reduces trajectory prediction error by 30% in high-interaction scenarios. Related recent work has also moved toward interpretable and causal machine learning for traffic safety, including causal forest-based incident duration analysis by Tahmid et al. (*16*) and causal mediation frameworks for roadway crash risk using connected vehicle data by Wang et al. (*17*). On the application side, Ribeiro et al. (*18*) demonstrated that LSTMs applied to V2X communication data can predict 96% of simulated intersection collisions with multi-second lead time, though a persistently high false positive rate highlights the difficulty of translating high recall into operationally viable warning systems — a tension this study confronts directly.

METHODOLOGY

Dataset and Scenarios

This study uses the INTERACTION Dataset v1.2 (*14*), a drone-recorded naturalistic driving dataset providing 10 Hz trajectory data with HD semantic maps at unsignalized intersections across the United States, Germany, and China. Two USA scenarios are used: DR_USA_Intersection_EP0, a T-junction with a single stop-controlled minor approach, and DR_USA_Intersection_MA, an all-way stop intersection with four controlled approaches. These two scenarios represent geometrically distinct intersection types, enabling an assessment of cross-scenario generalizability. A third scenario, DR_USA_Intersection_EP1, was preprocessed but excluded from modeling: feature extraction yielded only five usable events with no safe-class representation, making it unsuitable as a training or test source.

Preprocessing

Raw trajectory CSVs were loaded and validated against the HD map for each scenario. Minor-road candidate vehicles were identified by applying a stop-line proximity threshold of 2.0 m and a heading filter aligned to the stop-controlled approach direction. For EP0, a southbound heading filter (center $= -1.57$ rad, tolerance $= 1.10$ rad) isolated the single minor-road approach at the west junction;

for MA, all four approaches were treated as stop-controlled and processed independently, with east-west and north-south same-road pairs subsequently excluded to retain only crossing-approach conflict pairs. Figure 1 depicts the entire pipeline in the form of a flowchart.

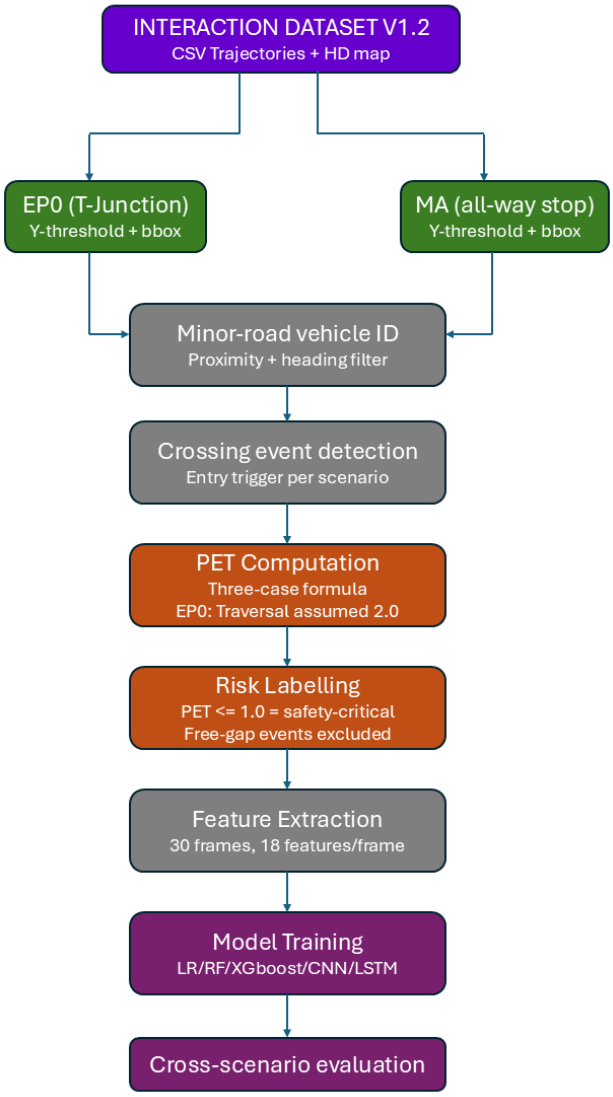

FIGURE 1 End-to-end pipeline from raw trajectory data to model evaluation.

The conflict zone – the spatial region where crossing paths intersect – was defined separately for each scenario and serves two distinct purposes: crossing detection and PET occupancy measurement. For EP0, crossing detection uses a y-coordinate threshold ($y < 990$) as the entry trigger, while PET computation uses a fixed bounding box ($x \in [983, 1010]$, $y \in [984, 990]$, approximately 42 m^2) to detect major-road vehicle occupancy. These definitions are intentionally decoupled: entry detection and conflict occupancy measurement serve different operational purposes. For MA, a central conflict box ($x \in [1002, 1036]$, $y \in [993, 1011]$, approximately 612 m^2) derived from the four stop-line centroids was used for both purposes. Figure 2 depicts conflict zone definitions for

EP0 (left) and MA (right), overlaid on stop-line geometry and vehicle trajectories.

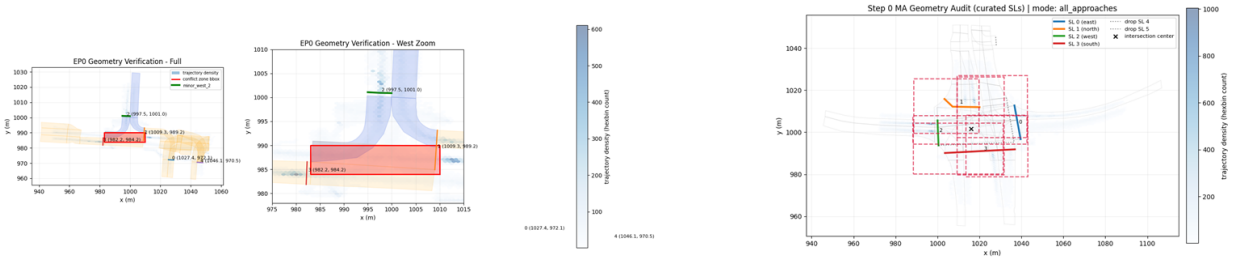

FIGURE 2 Conflict zone definitions for EP0 (left) and MA (right), overlaid on stop-line geometry and vehicle trajectories.

PET Computation and Risk Labeling

Post-Encroachment Time was computed for each crossing event using the conflict zone definitions above. Let $t_{\text{ego,entry}}$, $t_{\text{ego,exit}}$ denote the timestamps at which the minor-road vehicle enters and exits the conflict zone, and $t_{\text{conf,entry}}$, $t_{\text{conf,exit}}$ those of the conflicting major-road vehicle. PET is defined as:

$$\text{PET} = \begin{cases} t_{\text{ego,entry}} - t_{\text{conf,exit}} & \text{if conflict exits before ego enters} \\ 0 & \text{if occupancy windows overlap} \\ t_{\text{conf,entry}} - t_{\text{ego,exit}} & \text{if ego exits before conflict enters} \end{cases} \quad (1)$$

All 64 EP0 crossing events had clip-truncated exits due to the 4-second clip structure of the INTERACTION dataset. A fixed minor-road traversal time of 2.0 seconds was assumed to estimate $t_{\text{ego,exit}}$ for these events, and this limitation is acknowledged in the dataset constraints. For MA, clean exit timestamps were available and truncated events were excluded. Events with no conflicting vehicle present in the conflict zone window (free-gap events) were removed from both scenarios. Crossing events were then labeled using a threshold of PET $\leq$ 1.0 s as safety-critical (SC = 1) and PET $>$ 1.0 s as safe (SC = 0), following Hydén (*6*). This yielded 45 labeled events for EP0 (SC = 29, Safe = 16) and 978 for MA (SC = 858, Safe = 120), for a combined dataset of 1,023 events. Figure 3 demonstrates the PET distribution across EP0 and MA scenarios. The dashed vertical line at 1.0 s marks the safety-critical threshold.

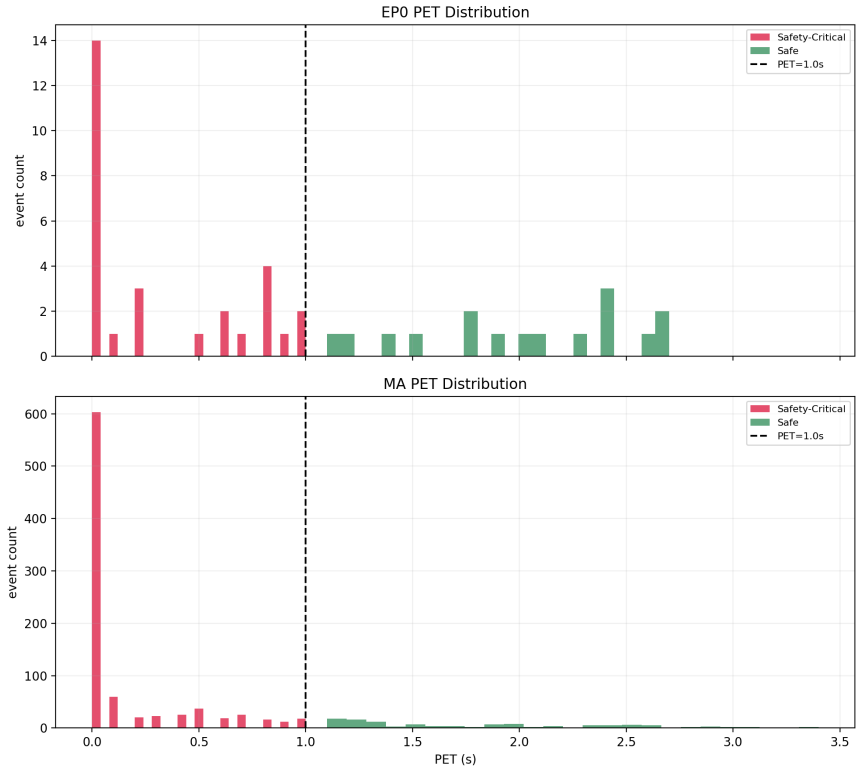

FIGURE 3 PET distribution across EP0 and MA scenarios.

Feature Engineering

For each labeled event, a fixed-length pre-entry observation window of 30 frames (3 seconds at 10 Hz) ending at the conflict zone entry timestamp was extracted. This window captures whatever approach behavior the driver exhibited prior to crossing – decelerating, waiting, or rolling through – without imposing any assumption about what that behavior looks like. Events with windows shorter than 30 frames were zero-padded at the start; 8 EP0 events required padding, while all MA events had complete windows.

Each frame contains 18 features organized in three groups. The 7 ego-vehicle features are:

$$x_{\text{rel}},\ y_{\text{rel}},\ v_x,\ v_y,\ \text{speed},\ a_x,\ a_y$$

The 7 conflict vehicle features follow the same form expressed relative to the ego frame. The 4 pairwise features are inter-vehicle distance, delta speed, delta heading, and time to zone. Ego coordinates are normalized to the origin at each frame:

$$x_{\text{rel}}^{(t)} = x^{(t)} - x^{(t)}, \quad y_{\text{rel}}^{(t)} = y^{(t)} - y^{(t_{\text{entry}})} \tag{2}$$

This position-invariant representation ensures the model learns from motion patterns rather than absolute location within the map. The conflict vehicle was identified as the major-road vehicle yielding the minimum PET within the conflict zone during the ego crossing window for EP0, and as the approach vehicle with minimum PET at the decision point for MA. The final feature tensors have shape $(N, 30, 18)$, where $N = 45$ for EP0 and $N = 978$ for MA.

Model Architectures and Training

Five classifiers were evaluated to span a range of complexity and temporal awareness. Logistic Regression (LR) served as the linear baseline, operating on flattened feature vectors. Random Forest (RF) and XGBoost extended this with ensemble tree methods capable of capturing non-linear feature interactions. A 1D Convolutional Neural Network (1D-CNN) applied temporal convolutions across the 30-frame window, and a Long Short-Term Memory network (LSTM) network with two stacked layers (`hidden_size=64`, `dropout=0.3` between layers) followed by an MLP head ($64 \rightarrow 32 \rightarrow 1$) modeled sequential dependencies across the pre-entry trajectory.

All models were trained on MA and evaluated on the held-out EP0 test set, assessing cross-scenario generalizability across geometrically distinct intersection types. MA events were split into train and validation sets using the first fold of `StratifiedGroupKFold` ($n\_splits = 5$, grouped by `case_id`), yielding 787 training and 191 validation events. EP0 serves as the sole held-out test set. Class imbalance (SC = 86.7% combined) was handled via `class_weight="balanced"` for LR and RF, `scale_pos_weight` for XGBoost, and `BCEWithLogitsLoss` with `pos_weight` for the deep learning models. Deep models were trained with the Adam optimizer (lr = 10^{-3}) and early stopping on validation AUC with patience of 10 epochs. Model selection prioritized AUC as the primary metric, with SC recall as a secondary criterion given the safety-critical application.

RESULTS

Dataset Summary

Table 1 summarizes the final labeled dataset across both scenarios. The combined dataset of 1,023 events is heavily imbalanced toward the safety-critical class (86.7%), driven primarily by MA where short accepted gaps in a dense all-way stop environment produce a high proportion of low-PET crossings. EP0 is more balanced by comparison, with 35.6% safe events, reflecting the less congested T-junction dynamics.

TABLE 1 Labeled dataset statistics by scenario.

Scenario	Total	SC	Safe	SC%
EP0 (T-junction)	45	29	16	64.4
MA (All-way stop)	978	858	120	87.7
Combined	**1,023**	**887**	**136**	**86.7**

Cross-Scenario Model Performance

All models were trained on MA and evaluated on the held-out EP0 test set (n=45). Deep models achieved validation AUC of approximately 0.97 on the MA validation split, indicating that all models learned the MA scenario well. The pronounced drop in performance on EP0 is therefore attributable to domain shift rather than underfitting, and constitutes the central finding of this study.

Table 2 reports full test metrics for all five models on EP0. The LSTM achieves the highest AUC (0.647) and maintains strong SC recall (0.862), making it the best overall model under a safety-critical framing. Random Forest achieves the highest raw SC recall (0.966) but at the cost of near-total safe class failure, predicting almost every event as safety-critical regardless of PET. Logistic Regression fails entirely, with AUC below chance, confirming that flattened tabular features without temporal structure carry insufficient discriminative signal for cross-scenario transfer.

TABLE 2 Model performance on EP0 held-out test set (n=45). Best value per metric in bold.

Model	Acc.	Prec.	Recall-SC	F1	AUC
Logistic Regression	0.422	0.615	0.276	0.381	0.418
Random Forest	0.622	0.636	**0.966**	0.767	0.544
XGBoost	0.600	0.628	0.931	0.750	0.535
1D-CNN	0.600	0.649	0.828	0.727	0.379
LSTM	**0.622**	**0.658**	0.862	**0.746**	**0.647**

Figure 4 provides a visual summary of performance across all metrics, with the dashed line marking the 0.9 recall target for safety-critical applications.

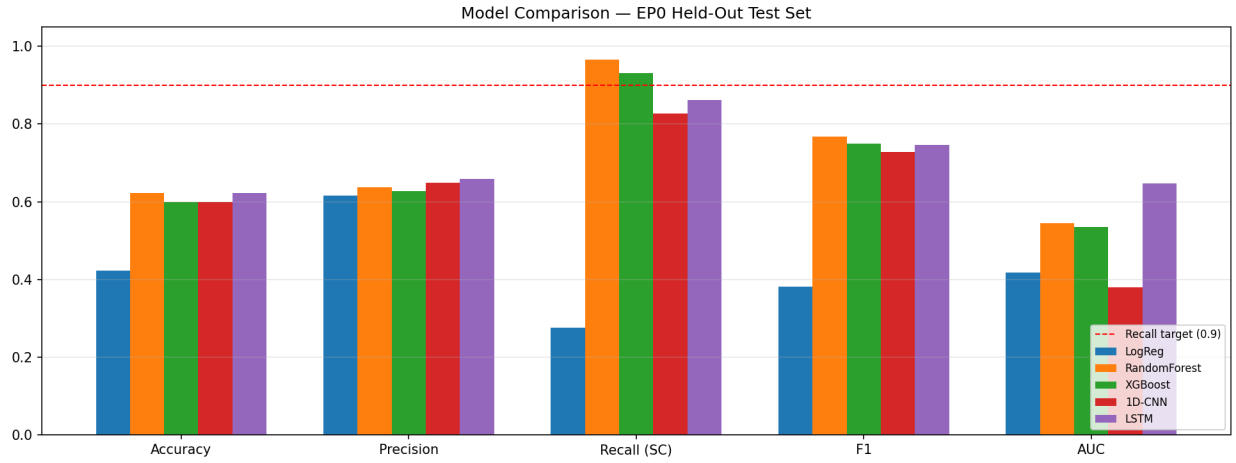

FIGURE 4 Model performance comparison across all metrics on the EP0 test set. Dashed line indicates the 0.9 SC recall target.

Feature Importance and Temporal Analysis

Figure 5 presents the Random Forest feature importance analysis aggregated across all 30 frames, alongside the temporal importance profile across the pre-entry window.

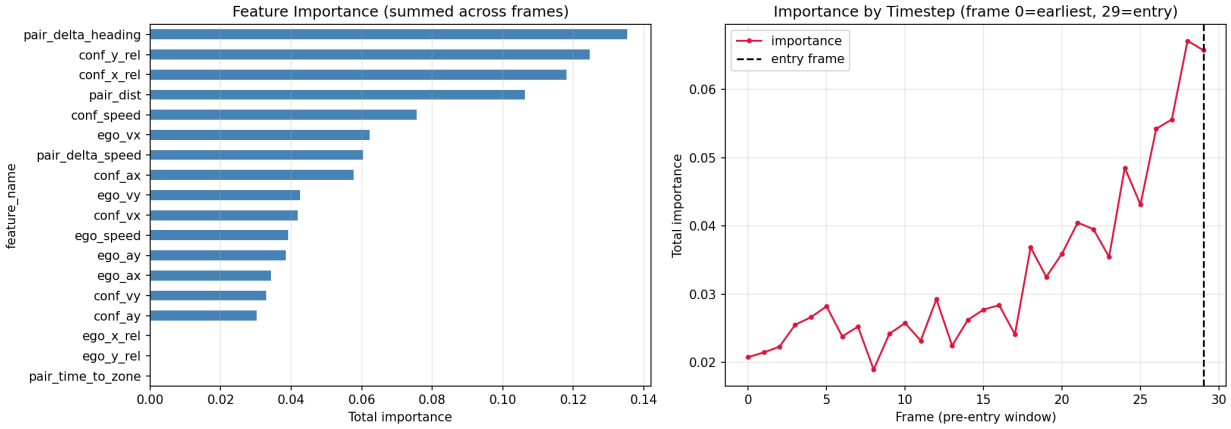

FIGURE 5 Left: aggregated RF feature importance by feature name. Right: total importance by timestep across the 30-frame pre-entry window (frame 0 = earliest, frame 29 = conflict zone entry).

Conflict geometry features dominate the importance ranking: `pair_delta_heading` (0.135), `conf_y_rel` (0.124), `conf_x_rel` (0.118), and `pair_dist` (0.106) account for over 48% of total importance collectively. Ego kinematic features contribute minimally, consistent with the EDA finding that ego entry speed carries no significant signal ($p = 0.89$). The temporal profile shows importance rising monotonically toward the entry frame, indicating that the last few frames before the vehicle commits to crossing are the most predictive of risk outcome. This aligns with the evidence accumulation account of gap acceptance decisions described by Zgonnikov et al. (*4*), where perceptual information closest to the decision point carries the greatest weight.

Error Analysis

Figure 6 shows the confusion matrices for the LSTM and Random Forest on the EP0 test set.

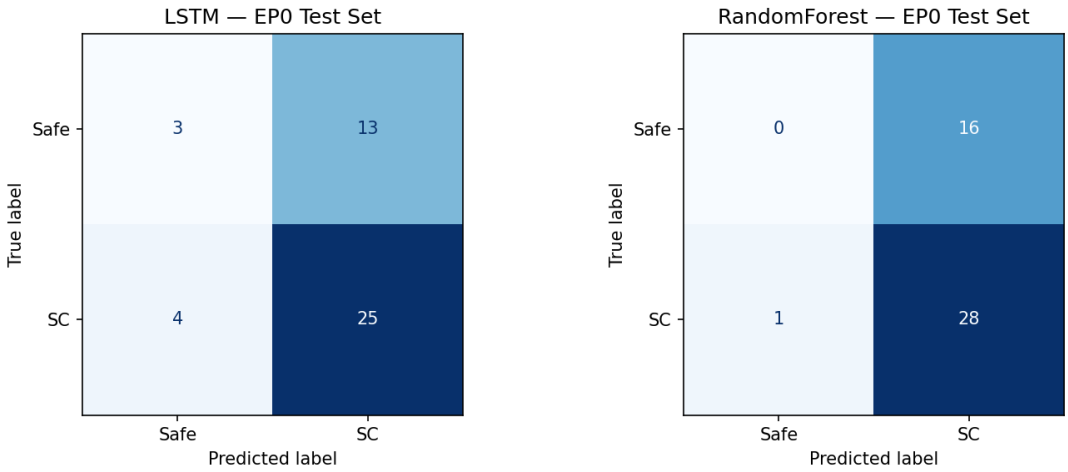

FIGURE 6 Confusion matrices for LSTM (left) and Random Forest (right) on the EP0 held-out test set.

The LSTM correctly identifies 25 of 29 SC events (4 false negatives) and 3 of 16 safe events (13 false positives). The 4 missed SC events have PET values of $\{0.0, 0.0, 0.0, 0.8\}$ seconds – three of which are overlap events (PET = 0), representing the most severe conflicts in the dataset. The model's failure on these events suggests that the most extreme conflicts may produce pre-entry trajectories indistinguishable from safe ones, consistent with the finding that SC recall degrades at the boundary of what trajectory data alone can resolve. The 13 false positives all have PET values above 1.0 seconds (range 1.1–2.7s), indicating the model over-predicts risk for geometrically ambiguous safe crossings. For a V2X warning system, this conservative failure mode – over-warning rather than under-warning – is preferable to the alternative.

DISCUSSION AND LIMITATIONS

The central finding – validation AUC of 0.97 on MA collapsing to 0.65 on EP0 – is not a modeling failure. It reflects a fundamental geometric mismatch: MA's symmetric all-way stop produces conflict dynamics that do not transfer to EP0's asymmetric minor-major road hierarchy. The implication for V2X deployment is direct: scenario-specific training data is a requirement, not an optimization.

The INTERACTION dataset's 4-second clip architecture creates two compounding problems for safety analysis. Short clips severely limit observable completed crossings – only 45 usable EP0 events survived after free-gap exclusion and approach filtering – and clip-truncated exits forced a fixed 2.0-second traversal assumption for all EP0 PET values, introducing irreducible label uncertainty into the test set. This is an architectural mismatch between the dataset's design intent (trajectory prediction benchmarking) and the demands of PET-based labeling. The resulting small test set makes all reported EP0 metrics high-variance estimates that should be interpreted with appropriate caution.

These limitations point to a concrete next step. The UCF CitySim (*19*) dataset offers continuous long-duration recordings at urban intersections, eliminating clip truncation, enabling clean PET computation, and yielding substantially more labeled events per scenario. The pipeline

developed here transfers directly. We regard trajectory-based pre-crossing risk classification as a promising direction, and the domain shift finding provides a testable hypothesis: whether geometry-aware or multi-scenario joint training strategies can produce classifiers that generalize across intersection types – a tractable problem with direct implications for smart intersection safety systems.

CONCLUSION

This study developed and evaluated a machine learning pipeline for classifying accepted gap acceptance events as safe or safety-critical using pre-decision trajectory data, grounded in PET-based risk labeling from naturalistic drone footage. Across five classifiers, the LSTM achieved the best cross-scenario performance (AUC = 0.647, SC recall = 0.862), with conflict vehicle geometry emerging as the dominant risk predictor and temporal importance rising toward the intersection entry point. The pronounced domain shift from MA to EP0 establishes that intersection geometry critically shapes the risk signal – a finding that directly informs how training data must be structured for real-world V2X deployment. With longer-duration datasets and geometry-aware training strategies, trajectory-based pre-crossing risk classification represents a viable and underexplored path toward proactive smart intersection safety systems.

ACKNOWLEDGMENTS

The authors thank the proprietors of the INTERACTION Dataset (*14*) for making the data available for this research.

AUTHOR CONTRIBUTIONS

This work, from conceptualization through experimentation and documentation, was conducted solely by the author.

FUNDING

None

REFERENCES

1. Federal Highway Administration, *About Intersection Safety*, 2024, accessed: 2026.
2. Li, Y., H. Hao, R. B. Gibbons, and A. Medina, Understanding gap acceptance behavior at unsignalized intersections using naturalistic driving study data. *Transportation research record*, Vol. 2675, No. 9, 2021, pp. 1345–1358.
3. Nagalla, R., P. Pothuganti, and D. S. Pawar, Analyzing gap acceptance behavior at unsignalized intersections using support vector machines, decision tree and random forests. *Procedia Computer Science*, Vol. 109, 2017, pp. 474–481.
4. Zgonnikov, A., D. Abbink, and G. Markkula, Should I stay or should I go? Cognitive modeling of left-turn gap acceptance decisions in human drivers. *Human factors*, Vol. 66, No. 5, 2024, pp. 1399–1413.
5. Schumann, J. F., J. Kober, and A. Zgonnikov, Benchmarking behavior prediction models in gap acceptance scenarios. *IEEE Transactions on Intelligent Vehicles*, Vol. 8, No. 3, 2023, pp. 2580–2591.
6. Hydén, C., The development of a method for traffic safety evaluation: The Swedish Traffic Conflicts Technique. *Bulletin Lund Institute of Technology, Department*, , No. 70, 1987.
7. Laureshyn, A., Å. Svensson, and C. Hydén, Evaluation of traffic safety, based on micro-level behavioural data: Theoretical framework and first implementation. *Accident Analysis & Prevention*, Vol. 42, No. 6, 2010, pp. 1637–1646.
8. Anowar, P., N. Haque, M. A. Raihan, and M. Hadiuzzaman, Trajectory-based real-time pedestrian crash prediction at intersections: A novel non-linear link function for block maxima led Bayesian GEV framework addressing heterogeneous traffic condition. *arXiv preprint arXiv:2510.12963*, 2025.
9. Jahan, I. A., M. Abdel-Aty, and Z. Islam, Impact of Sensing Modality on Surrogate Safety Metrics: A Radar–camera Fusion Framework for Intersection Monitoring. *Available as a preprint at SSRN 6486578*, 2026.
10. Jahan, I. A., A. S. Huq, M. K. Mahadi, I. A. Jamil, and M. Z. Shahriar, RoadSense: A Framework for Road Condition Monitoring using Sensors and Machine Learning. *IEEE Transactions on Intelligent Vehicles*, 2024, pp. 1–12.
11. Jahan, I. A., M. Abdel-Aty, and Z. Islam, Drone-Supervised Multi-Modal Sensor Fusion for Infrastructure-Based Vehicle Detection in Bird's Eye View. *IEEE Internet of Things Journal*, 2026.
12. Jahan, I. A., I. Arafat Jamil, M. S. Hossain Fahim, A. Sabiha Huq, F. Faisal, and M. M. Nishat, Accident Detection and Road Condition Monitoring Using Blackbox Module. In *2022 Thirteenth International Conference on Ubiquitous and Future Networks (ICUFN)*, 2022, pp. 247–252.
13. Krajewski, R., J. Bock, L. Kloeker, and L. Eckstein, The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems. In *2018 21st international conference on intelligent transportation systems (ITSC)*, IEEE, 2018, pp. 2118–2125.
14. Zhan, W., L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kummerle, H. Konigshof, C. Stiller, A. de La Fortelle, et al., Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps. *arXiv preprint arXiv:1910.03088*, 2019.

15. Diehl, F., T. Brunner, M. T. Le, and A. Knoll, Graph neural networks for modelling traffic participant interaction. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, IEEE, 2019, pp. 695–701.
16. Tahmid, M. M., Estimating Lane Control Signs (LCS) and Variable Speed Limit (VSL) Effects on Expressway Incident Duration: A Double Machine Learning Causal Forest Approach. *Available at SSRN 6199138*, 2025.
17. Wang, C., M. Abdel-Aty, S. Zhai, A. S. M. N. Uddin, and Z. Islam, From prediction to explanation: A machine learning and causal mediation framework for roadway crash risk with connected vehicle data. *Transportation Research Part C: Emerging Technologies*, Vol. 183, 2026, p. 105479.
18. Ribeiro, B., M. J. Nicolau, and A. Santos, Using machine learning on v2x communications data for vru collision prediction. *Sensors*, Vol. 23, No. 3, 2023, p. 1260.
19. Zheng, O., M. Abdel-Aty, L. Yue, A. Abdelraouf, Z. Wang, and N. Mahmoud, CitySim: A drone-based vehicle trajectory dataset for safety-oriented research and digital twins. *Transportation research record*, Vol. 2678, No. 4, 2024, pp. 606–621.