Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 2390–2400
https://doi.org/10.54941/ahfe1008464

Operationally Conditioned Transit Injury Risk: A National Machine Learning Analysis of Safety, Mobility, and Socioeconomic Exposure

Timothy Kangogo[1], Oladotun Adeola[2], Barbara Asamoah[3], Anthony Dontoh[2], Judith Mwakalonge[2], Gurcan Comert[4], Varghese Vaidyan[5], and Vijay Bendigeri[6]

[1]Department of Civil and Mechanical Engineering Technology, South Carolina State University, Orangeburg, SC 29117, USA
[2]Department of Engineering, South Carolina State University, Orangeburg, SC 29117, USA
[3]North Carolina A&T State University, Greensboro, NC 27411, USA
[4]Department of Computational Data Science and Engineering, North Carolina A&T State University, Greensboro, NC 27411, USA
[5]Department of Computer and Cyber Sciences, Dakota State University, Madison, SD 57042, USA
[6]Independent Researcher, Sacramento, CA, USA

ABSTRACT

Transit injury risk may vary with operating conditions and community context, but the usefulness of these factors for safety screening must be assessed against agencies' prior injury experience. We examined whether operations, socioeconomic conditions and ambient mobility improve injury prediction beyond ridership exposure and history, while distinguishing prediction from moderation of associations. National agency-mode-month records linked transit safety and service data to headquarters-county context. The broad sample comprised 84,721 observations from 2019–2025. Poisson regression, negative binomial regression and gradient-boosted trees used identical samples within each comparison, with 2019–2023 training, 2024 validation, retrospective evaluation in January–June 2025 and a reserved July–December 2025 evaluation after model settings were fixed. History reduced validation Poisson deviance by 37.5% beyond exposure, mode and calendar controls. A smoothed historical-rate benchmark remained competitive, with close calibration of reserved-period national totals but variation across subgroups. Socioeconomic variables improved boosted validation deviance by 0.0219; gains in both 2025 evaluations were positive but uncertain (0.0279 and 0.0177; 95% confidence intervals spanning zero). Headquarters-county mobility showed no reliable incremental improvement in either available period; earlier exploratory tests did not establish mobility–vulnerability moderation. A secondary bus analysis found no clear difference between collision and non-collision service-hour associations: the ratio of rate ratios for 10% greater hours was 0.985 (95% confidence interval 0.936–1.035). These findings support historical benchmarking and subgroup calibration before local safety review. They neither identify causal service effects nor establish the absence of mobility-related vulnerability.

Keywords: Transit safety, Injury prediction, Historical benchmarks, Exposure, Calibration, Socioeconomic context, Ambient mobility

Received September 26, 2026; Revised October 29, 2026; Accepted November 5, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

INTRODUCTION

Transit agencies need a defensible basis for prioritizing safety review. We test whether operational and community data improve injury prediction beyond ridership and prior injury history. Such predictions can direct attention to where injuries are expected to concentrate; they cannot identify individual vulnerability, behavior or intervention effects.

Moderation means that associations vary across conditions; it differs from improved prediction after accounting for history. Earlier exploratory tests, the temporal prediction comparison and a secondary bus injury-type analysis address these distinct questions.

Strathman, Wachana and Callas (2010) studied TriMet bus incidents; Ziedan and Brakewood (2020) examined U.S. light-rail and streetcar safety. Zhou et al. (2020) compared injury severity by incident type, and Lwanga, Mwanga and Mrema (2022) studied non-collision injuries. The incremental value of county context beyond history remains unresolved.

Exposure relationships vary by crash type (Qin, Ivan and Ravishanker, 2004), and specification affects interpretation (Lord and Mannering, 2010). We compare predictors on identical observations; secondary passenger-injury coefficients do not enter national all-person predictions.

DATA AND ANALYTICAL POPULATIONS

The unit is agency × mode × month, aggregating service types. National Transit Database Monthly Modal injuries (June 18, 2026 extract) include injuries to passengers and non-passengers; injuries per trip represent service-normalized burden, not passenger probability. Unlinked passenger trips (UPT), vehicle revenue hours (VRH) and vehicles operated in maximum service (VOMS) came from April 2026 ridership data (Federal Transit Administration, 2026a). June 18 Major/Non-Major extracts supplied bus classifications (Federal Transit Administration, 2026b). Monthly totals are not adjusted for historical reporting thresholds (Federal Transit Administration, n.d.).

Eligibility required annual Full Reporter status, an annual safety record and positive UPT, VRH and VOMS; 2024 status provisionally defined 2025 eligibility. Vanpool, Federal Railroad Administration modes, documented exceptions and non-Full/unknown labels were excluded. County recovery expanded coverage, but agreement between data sources does not establish complete injury reporting.

The eligible panel contained 97,024 observations, including 2018 history initialization. Complete operations and selected American Community Survey (ACS) variables yielded 84,721 modeling/evaluation rows for 2019–2025; complete mobility yielded 77,458 through June 2025 (Table 1). Comparisons use identical rows within each sample.

Table 1: Predictive samples by period.

Period	Broad Rows	Agencies	Injuries	Mobility Rows
2019–2023 training	60,434	517	95,081	59,620
2024 validation	12,225	493	22,853	11,949
Jan–Jun 2025 retrospective	6,027	483	10,722	5,889
Jul–Dec 2025 reserved	6,035	483	11,280	Unavailable

Broad-sample agencies and injuries are shown; agencies recur. H1: January–June; H2: July–December. Mobility training, validation and H1 total 77,458 rows. Bus analyses use a separate cohort.

ACS five-year estimates (U.S. Census Bureau, n.d.) describe headquarters counties, not service catchments, passenger residences or injury locations. Connecticut towns were linked to 2022 planning regions using Census relationships; other recoveries used documented addresses, with temporal proxies flagged. Partial Connecticut mobility aggregates were excluded.

Neighborhood Patterns block-group data (Advan Research, 2022), accessed through Dewey, supplied mobility through June 2025. Aggregation excluded invalid/nonfinite values, weighted valid dwell/distance values by positive device counts, preserved missing totals and computed stops per device from jointly observed records. Weighted block-group medians are not person-level medians; device sums are not unique residents. Analyst-selected screens required five records, complete positive counts, 80% valid device-weight coverage and 20 contributing device-weight units per specified measure. Set before evaluation, these screens do not establish representativeness.

We used complete cases, retaining reported zeros and excluding rows missing selected predictors. We did not impute predictors or replace missing outcomes with zeros. All 2025 rows used 2020–2024 ACS estimates. These describe the full five-year interval, and adjacent vintages overlap (U.S. Census Bureau, 2026).

METHODS

Predictor Blocks and Historical Benchmark

The sequence adds history, operations, ACS and mobility to exposure and controls. Logarithms are natural; linear models fix the log(UPT) coefficient at one. Controls are mode group (bus, demand response, ferry, rail), month indicators and numeric year. Agency identifiers, principal components, static density and current-outcome predictors were excluded.

History uses earlier agency-mode records, smoothed toward the earlier mode rate with study-selected prior exposure a = 1,000,000 trips:

$$r_i = (Y_{i,prior} + a\, r_{m,prior}) / (UPT_{i,prior} + a).$$

Here i is agency-mode, m is mode, $Y_{i,prior}$ is earlier injuries, $UPT_{i,prior}$ earlier trips and $r_{m,prior}$ the earlier mode rate with pooled fallback. Missing months add neither trips nor zeros. Evaluation histories end in December 2023. Predictions equal UPT × r_i; trips-only predictions use the pooled training rate. The empirical-Bayes-style own-history weight is $UPT_{i,prior}/(UPT_{i,prior} + a)$, with fixed a and a mode-rate reference rather than a fitted safety performance function (Hauer et al., 2002).

History comprises log(1 + smoothed injuries per million trips), log(1 + accumulated UPT), observed historical months and a no-history indicator. Operations are log(VRH) and log(VOMS).

The 26 ACS predictors are six commuting shares (driving alone, carpool, transit, walking, cycling and working from home); four household vehicle-availability shares (zero, one, two, three or more); poverty, Supplemental

Nutrition Assistance Program (SNAP) receipt, public assistance, unemployment, no internet, renter occupancy and rent burden ≥30%; median household income and age; shares aged 65+ and under 18; non-Hispanic Black, Hispanic and non-Hispanic Native Hawaiian/Pacific Islander shares; bachelor's degree or higher; and living alone. Shares are percentages of their table universes: workers aged 16+ for commuting; households for vehicle availability, SNAP, assistance, internet and living alone; the poverty universe; civilian labor force; occupied housing units for renting; renters with computable rent burden; population for age and race/ethnicity; and persons aged 25+ for education. Income is in dollars and median age is in years.

The seven mobility predictors are log(1 + device sum), log(1 + stop sum), stops per device, dwell minutes, log(1 + home distance in meters), log(1 + daytime distance in meters) and contributing block-group count.

Estimation and Evaluation

Poisson generalized linear models (GLMs) and negative binomial regression with quadratic variance (NB2) used log links and log-UPT offsets. NB2 variance is $\mu + \alpha\mu^2$, with expected count μ and training-estimated dispersion α. LightGBM gradient-boosted trees used Poisson loss on injury rates per million trips, weighted by exposure in millions; predictions were converted to counts (Ke et al., 2017). Mode/month encoding and standardization used training data only. Linear models omitted the three-or-more-vehicle reference share and training-constant columns.

Training covered 2019–2023. LightGBM used learning rate 0.03, 15 leaves, minimum leaf size 100, L2 regularization penalty 1 and at most 1,500 iterations, stopping after 100 without improved 2024 deviance. Thus 2024 is validation; January–June 2025 (H1) is retrospective. All 27 fits (three families × four broad-sample stages and three families × five mobility stages) and preprocessing were fixed before July–December 2025 (H2) evaluation. Twelve broad-sample models and two benchmarks were scored once without refitting or recalibration. H2 counts had been audited, but model performance remained reserved. We specified this comparison after exploration and before fitting it.

After H2 evaluation, but before calculating alternatives, we specified a standalone-benchmark sensitivity using a = 250,000, 500,000, 1,000,000, 2,000,000 and 4,000,000 trips. Historical inputs and evaluation rows stayed fixed. Fitted models retained a = 1,000,000, so this post hoc analysis tests benchmark sensitivity only.

Mean Poisson deviance was $D = 2n^{-1}\Sigma[y \log(y/\mu) - (y - \mu)]$, where y is observed injuries, μ predicted injuries and n the number of observations; 'the term $y \log(y/\mu)$ is defined as zero when $y = 0$. Lower is better; gain is earlier-stage minus later-stage deviance. Percentile 95% confidence intervals (CIs) used 2,000 paired whole-agency bootstraps, conditional on fitted models, excluding training/tuning uncertainty and unadjusted across families and secondary contrasts.

The primary contrast adds mobility after history, operations and ACS on identical complete-case rows, without statistical matching. Calibration compares observed/predicted totals (O/P); below one means overprediction. Groups use mode, thirds of agencies' mean annual training UPT and fixed

prediction bins. Screening measures injuries captured in the top 10% of agencies by predicted counts, rounded upward, versus ridership/history. It assesses later-period burden, not preventable excess risk or unseen-agency transfer.

Secondary Bus Passenger Injury Analysis

The secondary 2018–2024 cohort includes motor bus, commuter bus, bus rapid transit and trolleybus. Major rider/waiting/leaving and Non-Major customer injuries were classified as collision, non-collision, security or ambiguous, with security taking precedence. Reporting categories do not establish clinical severity, and Non-Major summaries cannot identify specific injury mechanisms. Twenty-one eligible bus months with inconsistent passenger totals across sources were excluded and then restored in a sensitivity analysis.

There were 41,276 agency-mode-months, 456 agencies and 553 units: 20,692 collision and 38,014 non-collision injuries; 2,845 security injuries and one ambiguous injury were excluded. The two outcomes were stacked in NB2 with a log-UPT offset, type-specific within-unit deviations and means of log VRH/log VOMS, and mode, month and year effects (Bell and Jones, 2015). No unit fixed intercepts were included. Agency-clustered covariance handles repeated/paired outcomes; decomposition does not remove confounding.

The principal contrast $\theta = \beta_{H,noncollision} - \beta_{H,collision}$ compares within-unit log-VRH coefficients. The ratio of rate ratios (RRR), $\exp[\theta \log(1.10)]$, represents 10% greater hours at fixed trips/fleet. Sample, estimator, dispersion and reporting sensitivities preceded fitting. Flagged zero-event years had ≥6 months, mean VOMS ≥100 and zero events, without assuming reporting failure. A free raw log-UPT sensitivity combines within/between associations; its coefficient contrast was exploratory.

Earlier Exploratory Analyses

Earlier NB2 moderation tests used 47,185 complete bus/rail rows before county recovery, with a log-UPT offset, spline main effects, calendar controls and agency-clustered covariance. Principal-component indices formed mobility–vulnerability, mode and service interaction blocks for total and passenger/patron injuries. Six tests used Holm (1979) adjustment, without recovered-cohort repetition. Shapley additive explanations (SHAP) described predictions, not moderation (Lundberg and Lee, 2017).

RESULTS

Historical Information and Incremental Context

In 2024, history reduced GLM deviance from 1.889 to 1.181: 37.5% beyond exposure/mode/calendar controls (gain 0.7075; 95% interval 0.3910–1.1488). LightGBM gained 0.7102 (0.4356–1.0643). The simple history benchmark achieved 1.172 versus the ACS GLM's 1.173 (Table 2). Its 50.7% reduction against trips-only (2.377) uses a different baseline. History supplied most of the improvement; these percentages describe prediction error, not injury

variance explained. In the post hoc comparison, LightGBM + ACS gained 0.0506 (95% CI 0.0205–0.0860) over the historical benchmark in 2024 and 0.0041 (–0.0645 to 0.1100) in H1; 2024 also guided early stopping.

Table 2: Broad-sample predictive performance.

Model and Cumulative Stage	2024 D	H1 2025 D	H2 2025 D	H2 O/P
Trips-only benchmark	2.377	2.274	2.439	0.953
Historical benchmark	1.172	1.226	1.226	1.016
Poisson GLM controls	1.889	1.878	1.921	0.807
Poisson GLM + history	1.181	1.263	1.259	0.911
Poisson GLM + operations	1.181	1.263	1.260	0.909
Poisson GLM + ACS	1.173	1.272	1.266	0.890
NB2 controls	2.796	3.066	3.272	0.496
NB2 + history	1.294	1.478	1.549	0.714
NB2 + operations	1.294	1.476	1.545	0.716
NB2 + ACS	1.333	1.513	1.588	0.712
LightGBM controls	1.859	1.786	1.826	0.937
LightGBM + history	1.149	1.242	1.223	0.925
LightGBM + operations	1.143	1.250	1.220	0.921
LightGBM + ACS	1.122	1.222	1.202	0.937

D: mean Poisson deviance; O/P: observed/predicted injuries. GLM: generalized linear model; NB2: negative binomial with quadratic variance; ACS: American Community Survey. Stages accumulate after trip exposure, mode, month and year controls. H1: January–June; H2: July–December. H2 excludes mobility.

Operations added little after history: GLM gain 0.0003 and LightGBM 0.0056, both CIs including zero. Broad-sample ACS improved LightGBM by 0.0219 (0.0037–0.0408), about 1.9%; GLM gained 0.0082 (–0.0065 to 0.0233), and NB2 worsened. The common-sample boosted ACS gain was 0.0135, with a CI spanning zero. In H1, the boosted ACS gain was 0.0279 (–0.0093 to 0.0735).

Mobility did not reliably improve prediction after history, operations and ACS in either available period (Table 3). LightGBM's 2024 gain was –0.0015 (–0.0139 to 0.0114). This is inconclusive evidence about the specified headquarters measures, not proof that mobility is irrelevant.

Table 3: Incremental mobility performance on common observations.

Model	Period	Gain in D	95% interval
Poisson GLM	2024	–0.0034	–0.0158 to 0.0079
Poisson GLM	H1 2025	–0.0145	–0.0350 to 0.0058
NB2	2024	0.0181	–0.0088 to 0.0554
NB2	H1 2025	0.0344	–0.0172 to 0.1171
LightGBM	2024	–0.0015	–0.0139 to 0.0114
LightGBM	H1 2025	–0.0097	–0.0346 to 0.0100

Gain in D (mean Poisson deviance): without minus with mobility, after history/operations/ACS; positive favors mobility. GLM: generalized linear model; NB2: negative binomial with quadratic variance; ACS: American Community Survey. H1: January–June. Rows: 11,949 (2024), 5,889 (H1). Agency-bootstrap intervals condition on fitted models; no July–December mobility.

Earlier moderation tests were inconclusive (unadjusted p = 0.207–0.715; all Holm-adjusted p = 1.000). They neither establish absent moderation nor replace recovered-cohort testing.

Reserved Evaluation and Calibration

In H2, the boosted ACS gain was 0.0177 (−0.0143 to 0.0452); GLM and NB2 gains were −0.0067 (−0.0300 to 0.0162) and −0.0430 (−0.1186 to 0.0093). Boosted ACS estimates were positive across periods but uncertain in both 2025 evaluations. The same agencies and fitted models recur, so these are not independent confirmations.

ACS LightGBM's paired H2 gain over the historical benchmark was 0.0241 (−0.0264 to 0.0773), leaving superiority unestablished. Against 11,280 observed injuries, history predicted 11,101 (O/P 1.016), ACS LightGBM 12,038 (0.937), GLM 12,675 (0.890) and NB2 15,840 (0.712). History was closer to H1/H2 national totals; in 2024 ACS LightGBM was closer (0.992 versus 1.073). NB2 overpredicted totals; the present analyses did not identify the cause. No evaluation-derived recalibration was applied.

National agreement concealed subgroup error (Figure 1). History underpredicted bus injuries and overpredicted rail injuries; ACS LightGBM was closer for bus, ferry and the small/medium size groups. Neither model was uniformly best calibrated. H2 top-decile capture was 78.0%, 77.8% and 75.9% for history, ACS LightGBM and ridership; these descriptive differences do not establish a ranking advantage.

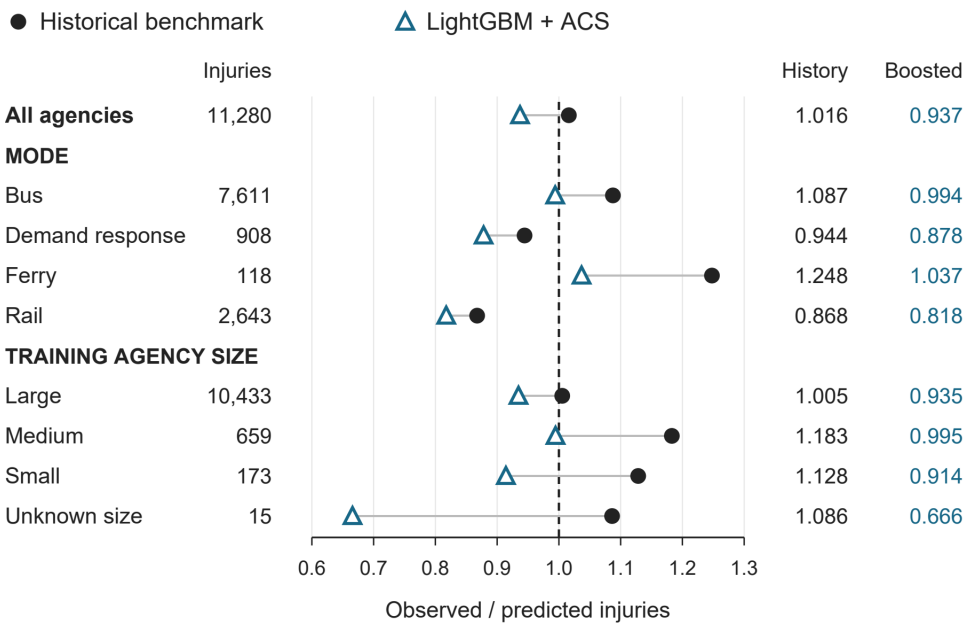


Figure 1: July–December 2025 calibration by mode and training agency-size group. Ratios below 1 indicate overprediction; above 1, underprediction. Values are descriptive, without uncertainty intervals. Unknown size includes 90 rows with 15 injuries; rows in different groupings overlap.

Across post hoc smoothing weights, benchmark deviance ranged from 1.170 to 1.200 in 2024, 1.224 to 1.246 in H1 and 1.222 to 1.249 in H2; H2 O/P ranged from 1.012 to 1.025. The original weight was retained.

Secondary Bus Contrast and Exposure Sensitivity

The principal difference was –0.163 (95% CI –0.691 to 0.365; p = 0.546), giving RRR 0.985 (0.936–1.035): 6.4% lower to 3.5% higher per 10% greater hours. This contrast does not establish different service-hour associations or equivalence; no equivalence margin was specified.

Table 4: Bus service-hour contrasts and sensitivity analyses.

Run	Months	RRR	95% CI
Principal NB2	41,276	0.985	0.936 to 1.035
Exclude 2020–2021	29,478	0.982	0.915 to 1.055
Restore source-discrepant months	41,297	0.984	0.936 to 1.035
Extend through June 2025	44,208	0.989	0.940 to 1.041
Poisson unit-by-type fixed effects	29,778	0.968	0.914 to 1.025
Estimate log-UPT coefficients	41,276	0.956	0.907 to 1.008
Exclude flagged zero-event years	41,156	0.983	0.935 to 1.034
Major-only non-collision outcome	41,276	0.924	0.836 to 1.021
Type-specific NB2 dispersion	41,276	0.983	0.936 to 1.033
Pooled passenger/patron outcome	41,276	N/A	Pooled slope reported below

RRR: ratio of injury rate ratios for 10% greater service hours, non-collision relative to collision. CI: confidence interval, clustered by agency. NB2: negative binomial with quadratic variance; UPT: unlinked passenger trips. Months are agency-mode-months. Fixed effects omit 357 all-zero unit-type cells; months count their retained union. N/A: no pooled between-type contrast; 2025 eligibility is provisional.

All nine contrast intervals included one (Table 4). The Major-only estimate, 0.924 (0.836–1.021), used 1,621 non-collision injuries, leaving reporting-threshold explanations unresolved. The pooled slope was –0.468 (–0.737 to –0.199).

Exploratory raw log-UPT coefficients were 0.035 for collision (CI –0.206 to 0.276) and 0.312 for non-collision (0.140–0.484); their difference was 0.276 (0.03–0.52; p = 0.026). A joint test rejected coefficients of one. These mixed within/between associations do not test the national all-person exposure assumption.

DISCUSSION

Prior injury history is a demanding baseline for new contextual predictors. We interpret it as potentially summarizing hazards, operating environments and reporting practices that this predictive design cannot distinguish. The small, uncertain gains from added context therefore concern incremental

predictive value; they do not show that operations or community conditions are unimportant to safety.

The secondary bus contrast remains inconclusive and does not justify changing exposure denominators (Qin et al., 2004; Lord and Mannering, 2010).

Headquarters measures may misrepresent routes or passenger origins. Research on SafeGraph documents spatial and temporal variation in mobile-location sampling bias (Li et al., 2024); the magnitude and direction of error in our mobility measures remain unknown. H2 used carried-forward ACS and lacked mobility. Findings apply to these measures, not all contextual relationships.

For safety review, compare added predictors with trips and history on identical observations, then check calibration by mode and agency size. Figure 1 shows why national agreement alone is insufficient: opposing subgroup errors can misdirect review priorities. High predicted counts can prioritize local investigation, but cannot justify automatic service reductions or unsafe-agency labels. Burden is neither preventable harm nor an equity criterion; mechanisms such as boarding or on-board movement require local event evidence.

LIMITATIONS AND FURTHER RESEARCH

Same-period service/ACS may be unavailable when decisions are made, limiting forecasting. Most units have histories. An earlier, separately fitted boosted model overpredicted held-out-agency totals by 31.4% in a combined agency/time split; that result does not evaluate the present models' transferability. Reporting thresholds, provisional eligibility, headquarters proxies, complete-case selection and shared counties may affect results. ACS sampling and training uncertainty were not propagated.

Approximately 96% of bus non-collision injuries are Non-Major, limiting mechanism interpretation and reporting comparability. National refits with free log-UPT or no linear year trend, pandemic/geography/zero-event exclusions and mobility-preparation variants remain incomplete; bus sensitivities do not establish national robustness. Updating histories beyond December 2023 requires equal treatment of benchmarks and fitted models. Within/between trip identification and independently evaluated recalibration also require further work.

CONCLUSION

Historical information provides a strong baseline, with close national calibration but subgroup errors. Added predictors should demonstrate improved prediction and acceptable calibration before guiding review priorities. Neither ACS in either 2025 evaluation nor mobility where available showed reliable gains; the secondary bus contrast was inconclusive. Evidence is strongest for temporal screening with established histories; unseen-agency generalization remains unestablished.

ACKNOWLEDGMENT

We thank our institutions and colleagues for their support.

REFERENCES

Advan Research. (2022). Foot Traffic / Neighborhood Patterns [dataset]. Dewey Data. doi: 10.82551/MS2A-2E59. https://doi.org/10.82551/MS2A-2E59

Bell, A., and Jones, K. (2015). Explaining fixed effects: Random effects modeling of time-series cross-sectional and panel data. Political Science Research and Methods, 3(1), 133–153. doi: 10.1017/psrm.2014.7.

Federal Transit Administration. (2026a). Complete Monthly Ridership (with adjustments and estimates) [dataset]. April 2026 workbook. https://www.transit.dot.gov/ntd/data-product/monthly-module-adjusted-data-release

Federal Transit Administration. (2026b). Major Safety Events and Non-Major Safety Events [datasets]. Study extracts dated June 18, 2026. https://data.transportation.gov/Public-Transit/Major-Safety-Events/9ivb-8ae9; https://data.transportation.gov/Public-Transit/Non-Major-Safety-Events/urir-txqm

Federal Transit Administration. (n.d.). Monthly Modal Time Series (Safety and Service). Archived study extract dated June 18, 2026. https://www.transit.dot.gov/ntd/data-product/monthly-modal-time-series-safety-and-service

Hauer, E., Harwood, D. W., Council, F. M., and Griffith, M. S. (2002). Estimating safety by the empirical Bayes method: A tutorial. Transportation Research Record, 1784(1), 126–131. doi: 10.3141/1784-16.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. Scandinavian Journal of Statistics, 6(2), 65–70. https://www.jstor.org/stable/4615733

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. Advances in Neural Information Processing Systems, 30. https://papers.nips.cc/paper_files/paper/2017/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html

Li, Z., Ning, H., Jing, F., and Lessani, M. N. (2024). Understanding the bias of mobile location data across spatial scales and over time: A comprehensive analysis of SafeGraph data in the United States. PLOS ONE, 19(1), e0294430. doi: 10.1371/journal.pone.0294430

Lord, D., and Mannering, F. (2010). The statistical analysis of crash-frequency data: A review and assessment of methodological alternatives. Transportation Research Part A: Policy and Practice, 44(5), 291–305. doi: 10.1016/j.tra.2010.02.001

Lundberg, S. M., and Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30. https://papers.nips.cc/paper_files/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html

Lwanga, A., Mwanga, H. H., and Mrema, E. J. (2022). Prevalence and risk factors for non-collision injuries among bus commuters in Dar es Salaam, Tanzania. BMC Public Health, 22, 963. doi: 10.1186/s12889-022-13284-9.

Qin, X., Ivan, J. N., and Ravishanker, N. (2004). Selecting exposure measures in crash rate prediction for two-lane highway segments. Accident Analysis & Prevention, 36(2), 183–191. doi: 10.1016/S0001-4575(02)00148-3.

Strathman, J. G., Wachana, P., and Callas, S. (2010). Analysis of bus collision and non-collision incidents using transit ITS and other archived operations data. Journal of Safety Research, 41(2), 137–144. doi: 10.1016/j.jsr.2010.02.003.

U.S. Census Bureau. (2026). Comparing ACS Data. Revised January 23, 2026. https://www.census.gov/programs-surveys/acs/guidance/comparing-acs-data.html

U.S. Census Bureau. (n.d.). American Community Survey 5-year estimates, detailed tables, 2015–2019 through 2020–2024 [datasets]. https://www.census.gov/data/developers/data-sets/acs-5year.html

Zhou, H., Yuan, C., Dong, N., Wong, S. C., and Xu, P. (2020). Severity of passenger injuries on public buses: A comparative analysis of collision injuries and non-collision injuries. Journal of Safety Research, 74, 55–69. doi: 10.1016/j.jsr.2020.04.003

Ziedan, A., and Brakewood, C. (2020). Longitudinal analysis of light rail and streetcar safety in the United States. Transportation Research Record, 2674(9), 83–95. doi: 10.1177/0361198120927004