Journal of High School Science

Crashmap360: a data-driven infrastructure recommendation system for Washington, D.C.

Kundra S

Submitted: December 29, 2025, Revised: version 1, May 25, 2026, version 2, June 10, 2026, version 3, June 27, 2026, version 4, July 9, 2026
Accepted: July 9, 2026

Abstract

Traffic safety remains a persistent challenge in Washington, D.C., despite initiatives such as Vision Zero and The Lab @ D.C. Rather than functioning as a general crash-prediction model, CrashMap360 is a decision-support framework that identifies candidate locations for infrastructure review by integrating crash severity, infrastructure proximity, traffic exposure, and spatial clustering. The framework was applied to three crash categories: nighttime crashes, speeding-related crashes, and cyclist-involved crashes. Within each category, crashes were filtered by type, distances to the nearest relevant safety infrastructure were calculated, and crashes beyond predefined infrastructure-distance thresholds were clustered and prioritized using crash frequency, severity-based metrics, and exposure-normalized risk. The robustness of the framework was evaluated through comparisons with capped-severity rankings, crash-frequency rankings, exposure-normalized rankings, DBSCAN clustering, the District Department of Transportation's High Injury Network, and threshold sensitivity analyses. The results identified candidate locations for engineering review across all three crash categories and showed that hotspot prioritization depends on the ranking methodology, with exposure normalization and severity weighting producing category-specific changes in prioritization while threshold perturbations had relatively little effect on the ranking of locations. These findings suggest that CrashMap360 provides a robust and flexible framework for systematically prioritizing candidate locations for engineering review while complementing existing transportation safety analyses.

Keywords

Traffic safety, Road safety, Crash hotspot analysis, Geographic Information Systems (GIS), Spatial clustering, Infrastructure prioritization, Vision Zero, Traffic crash analysis, Transportation planning, Decision support systems

Sachin Kundra, Sidwell Friends School, 3825 Wisconsin Ave NW, Washington, DC 20016, USA.
sdkundra@icloud.com

https://doi.org/10.64336/001c.165268

1. Introduction

Car accidents are one of the most common causes of death in the United States of America. In 2024, an estimated 40,000 lives were lost to traffic-related incidents (1). For teens between the ages of 13 to 19, car accidents are one of the leading causes of death (2). Motor vehicle crashes kill more children than influenza, COVID-19, or cancer. Unlike these diseases, however, many crash-related injuries and fatalities are preventable through infrastructure improvements and evidence-based transportation policies. Municipalities across the United States have implemented a wide range of interventions to reduce traffic fatalities, yet many jurisdictions continue to experience persistently high crash and fatality rates. Motor vehicle crashes were the leading cause of death among children aged 1–17 before 2020 and are now the second leading cause, indicating measurable progress. In 2023, 1,019 children died in motor vehicle crashes—the fewest since 2014, when 1,073 children were killed (3). Despite this improvement, motor vehicle crashes remain one of the leading causes of death among children and adolescents in the United States.

The D.C. government is making an effort. For example, Mayor Muriel Bowser adopted a project known as Vision Zero that focuses on reducing fatalities to zero in D.C. from traffic-related accidents. However, Vision Zero does not provide the crash-category-specific, infrastructure-linked ranking system developed in this study. Despite Vision Zero initiatives, severe injuries and traffic fatalities in the District have continued to increase (4).

The District of Columbia also evaluated a behavioral intervention to reduce dangerous driving. The Lab @ DC, an applied research and design team within the D.C. government, identified approximately 30,000 drivers at elevated risk of future crashes and notified them through text messages and mailed letters. However, these personalized risk notifications did not reduce subsequent crash involvement. The investigators suggested that the intervention's limited effectiveness may have resulted from the absence of meaningful consequences for risky driving. They also found that high-risk drivers were disproportionately from underserved and low-income communities, raising concerns that punitive measures could impose additional financial burdens on already vulnerable populations (5, 6).

The findings from The Lab @ DC suggest that interventions targeting roadway infrastructure may warrant greater emphasis than those targeting individual driver behavior alone. Because high-risk drivers were disproportionately from underserved and low-income communities, infrastructure-based strategies have the potential to improve safety without imposing additional financial burdens on vulnerable populations. Motivated by this perspective, CrashMap360 is a quantitative framework that integrates Washington, D.C. crash and infrastructure data to identify and prioritize locations for engineering review and potential safety interventions.

1.1 Literature review

Spatial analysis is often used in traffic safety to determine locations where crashes are highly concentrated. Strauss and Lentz demonstrated that this crash clustering can appear on different spatial scales. These scales can range from intersections to corridors and larger areas, and are important in determining what specific intervention may be necessary (7). Hovenden and Liu found that crash patterns varied substantially across crash categories in a study of metropolitan Melbourne, Australia (8). Collectively, these studies demonstrate that different crash categories exhibit distinct spatial patterns, supporting CrashMap360's use of category-specific hotspot analysis rather than a single aggregated analysis of all crashes.

When identifying crash hotspots, raw crash counts alone fail to account for differences in crash severity. Alam et al. incorporated both spatial patterns and crash severity in their analysis of traffic hotspots in Ohio, supporting the inclusion of severity as a criterion for hotspot prioritization (9). Similarly, crash frequency can be influenced by traffic exposure, as roads with higher traffic volumes inherently present more opportunities for collisions. Mansfield et al. addressed this issue by incorporating traffic exposure into models of pedestrian fatality risk (10). Consistent with these approaches, CrashMap360 integrates both crash severity and traffic volume to prioritize locations with the greatest potential for safety improvement.

Roadway interventions should still be interpreted carefully, as a review by Staton et al. found that evidence for road-traffic injury prevention interventions in low- and middle-income countries was limited, while legislation and enforcement often had stronger evidentiary support (11). The findings from the review suggest that infrastructure additions should not be treated as automatically effective in every location and that non-infrastructure-based approaches may be more appropriate in some contexts.

Research shows that clustering, severity, exposure, and infrastructure context are all important, but these different approaches can be incomplete when used alone. The literature review suggests that while these individual approaches are already well established, relatively few studies have combined these approaches into one infrastructure-specific prioritization framework. Crashmap360 incorporates all of these components, and contributes a severity- and exposure-aware prioritization framework that transforms crash data into a ranked set of candidate locations for safety review. Infrastructure proximity is used as a screening measure rather than causal evidence. This approach is consistent with Wang et al.'s built-environment analysis of crash hotspots in Kigali, which identified associations between hotspot characteristics and crash or mortality risk but did not claim causal relationships (12). Although the analysis focuses on Washington, D.C., the framework can be applied to other cities with comparable crash, infrastructure, and traffic-volume data.

2. Materials and Methods

Police accident reports were unavailable due to privacy reasons and a shortage of workers to honor the Freedom of Information Act (FOIA). Public information from a source termed Open Data DC was used instead (13). This resource contains maps and coordinates of infrastructure in the city that are essential to understanding the layout of the area. Open Data DC also contains the crash data used in the program. The data points useful to the project are shown in Table 1.

Table 1. Features of Open Data DC used in the study.

Feature	Description
REPORTDATE	The date and time that the car accident occurred.
MAR_SCORE	MAR score is a confidence rating that the Metropolitan Police Department (MPD) sets on their certainty of the accident; the scale ranges from 0 to 200.
LONGITUDE	The lateral coordinate for where a car crash happened.
LATITUDE	The vertical coordinate for where a car crash happened.
MAJORINJURIES_BICYCLIST	The number of cyclists involved in an accident who had major injuries.
MINORINJURIES_BICYCLIST	The number of cyclists involved in an accident who had minor injuries.
FATAL_BICYCLIST	The number of cyclists involved in an accident who died.
MAJORINJURIES_DRIVER and MAJORINJURIES_PASSENGER	The number of drivers and passengers involved in an accident who had major injuries.
MINORINJURIES_DRIVER and MINORINJURIES_PASSENGER	The number of drivers and passengers involved in an accident who had minor injuries.
FATAL_DRIVER and FATAL_PASSENGER	The number of drivers and passengers involved in an accident who died.
MAJORINJURIES_PEDESTRIAN	The number of pedestrians involved in an accident who had major injuries.
MINORINJURIES_PEDESTRIAN	The number of pedestrians involved in an accident who had minor injuries.
FATAL_PEDESTRIAN	The number of pedestrians involved in an accident who died.
SPEEDING_INVOLVED	The number of cars that are speeding and involved in the accident.

A common analytical framework was applied across the three crash categories: nighttime, speeding-related, and cyclist crashes. Pedestrian crashes were excluded because the available infrastructure layers did not permit reliable differentiation between crosswalks and sidewalks, reducing the specificity of proximity-based analyses. For each crash category, the crash dataset was first filtered by date, retaining crashes that occurred from January 1, 2020 through April 30, 2025. This timeframe captures changes in travel behavior associated with the COVID-19 pandemic and the expansion of app-based delivery services while remaining recent enough to reflect the current roadway infrastructure. The dataset is sufficiently long to provide an adequate sample size without including infrastructure that may have since changed. Finally, crashes with a MAR_SCORE below 100 were excluded to improve data quality and ensure that only records meeting the District's reporting threshold were analyzed.

After filtering, the distance from each crash to the nearest infrastructure feature intended to mitigate that crash type was calculated. Crashes located beyond a predefined distance threshold were retained, clustered, and ranked according to their severity scores. To account for differences in traffic exposure, exposure-normalized rankings were generated and compared with rankings based on raw crash severity. Finally, the resulting hotspot rankings were compared with those produced by the density-based clustering algorithm DBSCAN and with the District Department of Transportation's (DDOT) High Injury Network to evaluate methodological agreement and identify differences in hotspot prioritization.

2.1 Severity scoring

A severity score was constructed to quantify frequency and injury level in certain areas. An alternative system that accomplishes this task is the Equivalent Property Damage Only (EPDO) method, which weights crashes based on their estimated economic costs. In Sharafeldin et al.'s study, EPDO was used to combine frequency and severity into one measure (14). In Crashmap 360, the severity score method was used instead of EPDO so as to increase interpretability and to preserve the model's focus of identifying locations with a focus on the number of fatalities and major injuries rather than a high economic crash-cost estimation.

Cluster severity was calculated as the sum of the severity scores for all crashes within a cluster. Each crash received a severity score of [(7 × fatalities) + (4 × major injuries) + (1 × minor injuries)]. This weighting scheme emphasizes fatalities and major injuries because the primary objective of CrashMap360 is to prioritize locations with the greatest potential to reduce deaths and life-altering injuries. Minor injuries were assigned a lower weight but were retained because a concentration of minor-injury crashes may still indicate an unsafe location requiring infrastructure review. Fatalities were weighted more heavily than minor injuries while remaining comparable to multiple major injuries, balancing the prioritization of fatal crashes with clusters exhibiting repeated serious injuries. As with any severity-weighting framework, these values reflect the objectives of the study rather than universal standards, and alternative weightings may be appropriate for different policy or research priorities.

The calculation of this severity metric leads to the possibility of singular crashes dominating an entire cluster, as a cluster with a singular accident with many fatalities would be weighted much higher than a cluster with more accidents but fewer fatal crashes. Thus, a second severity sum was calculated, called the capped severity sum. For this additional calculation, singular crash severity values are capped at the highest degree of injury in that specific crash. For example, if there are multiple major and minor injuries within a crash, the crash would be capped at one major injury such that the crash would not dominate a cluster. To assess the robustness of each severity metric, the two severity frameworks were calculated alongside one another for each cluster. For all three crash types, the two tables of final ranked clusters were each ranked by severity sum and capped severity sum, and were then compared using Spearman rank correlation coefficient, an analysis that measures the similarities of rankings across two tables with the same contents but different orders (15). The Spearman rank correlation coefficient is computed by $\rho = 1 - \frac{6\sum_{i=1}^{n} d_i^2}{n(n^2-1)}$ (Eq.1), where ρ = Spearman's rank correlation coefficient, d_i = difference between the paired ranks of observation I, n = number of paired observations and $\sum_{i=1}^{n} d_i^2$ = sum of the squared rank differences across all observations. Values of ρ range from

−1 (perfect inverse monotonic association) to +1 (perfect monotonic association), with 0 indicating no monotonic relationship. Further analysis compared the 10 highest-severity clusters from each ranking using the Top-10 Overlap metric, defined as the percentage of clusters appearing in the top 10 of both rankings.

In addition to comparing the severity-sum and capped-severity-sum rankings, crash frequency was also evaluated. Frequency was defined as the total number of crashes within each cluster, and clusters were ranked accordingly in a separate table. Spearman's rank correlation coefficient and the Top-10 Overlap metric were then used to compare the frequency-based rankings with the severity-based rankings.

To address the potential conflation of numerous low-severity crashes with fewer high-severity crashes in the severity-sum metric, three additional prioritization metrics were evaluated alongside the capped severity sum and crash frequency. First, the average severity per crash was calculated as the cluster severity sum divided by the total number of crashes in the cluster. Second, the severe crash percentage was calculated as the proportion of crashes within a cluster that resulted in either a major injury or a fatality. Third, the fatal crash percentage was calculated as the proportion of crashes within a cluster that resulted in at least one fatality.

2.2 Night accidents

For nighttime crashes, the associated infrastructure was streetlights, with locations obtained from Open Data DC (16). Clusters of nighttime crashes occurring far from existing streetlights were identified as candidate locations for further lighting review. Because Washington, D.C. has 72,029 streetlights, a preliminary analysis was conducted to determine whether a sufficient number of crashes occurred beyond the effective illumination range of existing lighting. Streetlights typically illuminate an area within approximately 30–50 m (17). Using a 30 m threshold, 91.64% (15,611/17,035) of nighttime crashes occurred within range, leaving 8.36% (1,424/17,035) beyond the threshold. Increasing the threshold to 50 m reduced this proportion to 4.51% (769/17,035). Although most nighttime crashes occurred near existing streetlights, the 30 m threshold retained 1,424 crashes for clustering, providing a sufficient sample to identify locations where lighting conditions may warrant further engineering review.

Crash and streetlight coordinates were projected into a planar coordinate reference system (meters), allowing distances to be calculated using the Euclidean distance metric $d=\sqrt{(x_2-x_1)^2+(y_2-y_1)^2}$ (Eq. 2), where d = Euclidean distance, x_1,y_1 = coordinates of the crash and x_2,y_2 = coordinates of the nearest streetlight. Because the study is confined to Washington, D.C., distortion from the Earth's curvature is negligible, making Euclidean distance appropriate. To efficiently identify the nearest streetlight for each crash, a BallTree nearest-neighbor spatial index was used rather than a brute-force search. A brute-force approach would require approximately 1.22 billion distance calculations (17,000 crashes × 72,000 streetlights), whereas the BallTree recursively partitions the streetlight locations into spatial regions, substantially reducing the number of distance computations. Nighttime crashes occurring more than 30 m from the nearest streetlight were retained for clustering, while crashes within 30 m of an existing streetlight were excluded. This filtering reduced the dataset to 1,424 crashes, after which pairwise distance calculations (approximately 2.03 million comparisons) were used for clustering.

Hierarchical complete-linkage clustering was then applied using a 30 m linkage threshold. Complete linkage requires every pair of crashes within a cluster to be separated by no more than the threshold distance, preventing elongated chains of loosely connected crashes. The 30 m threshold was chosen to match the approximate effective illumination distance of a streetlight, ensuring that each identified cluster represented a localized area potentially underserved by existing lighting infrastructure.

2.3 Speeding

For speeding-related crashes, the locations of speed cameras, speed humps, and crashes were obtained from Open Data DC (18, 19). The same filtering procedures described previously were applied, except that only crashes involving speeding were retained and the nighttime restriction was removed. After filtering, the dataset contained 1,210 speeding-related crashes and 589 traffic-calming devices (speed cameras and speed humps). Because the dataset was relatively small, the distance from each crash to its nearest device was calculated using a brute-force nearest-neighbor search, requiring approximately 713,000 pairwise Euclidean distance calculations.

Unlike streetlights, speed cameras and speed humps influence traffic only along the roadway on which they are installed. Consequently, geographic proximity alone may not reflect the effectiveness of these devices. For example, a speed hump may reduce speeding on one street but have little or no effect on an adjacent perpendicular street despite their close spatial proximity.

To account for the roadway-specific influence of speed cameras and speed humps, crashes and traffic-calming devices were assigned to their nearest roadway segment using the District Department of Transportation (DDOT) roadway centerline dataset (20). This assignment reduced erroneous associations between crashes and devices located on nearby but different streets. Crashes greater than 20 m from a roadway centerline were removed. Distances were then calculated only between crashes and devices assigned to the same roadway segment.

A proximity threshold of 59.4 m, corresponding to the median DDOT roadway-segment length, was used to represent the approximate distance over which a speed camera or speed hump could influence driver behavior along a street segment. Of the 1,097 speeding-related crashes analyzed, 54 (4.92%) occurred within 59.4 m of an existing device, whereas 1,043 (95.08%) occurred beyond this threshold. This distribution suggests that most speeding-related crashes occur on roadway segments without existing speed-control devices, likely reflecting the relatively small number of such devices compared with the total number of roadway segments in the city.

The remaining crashes were clustered using hierarchical complete-linkage clustering with a 59.4 m linkage threshold. Applying complete linkage within individual roadway segments prevented clusters from spanning multiple streets while identifying localized groups of speeding-related crashes at the scale of a typical city block.

2.4 Bike lanes

The cyclist-crash analysis followed the same general framework as the nighttime and speeding analyses. For each cyclist-involved crash, the distance to the nearest bike lane was calculated. Crashes beyond a predefined distance threshold were retained, clustered, and ranked according to their severity sums.

After filtering and integrating crash and bicycle infrastructure data from Open Data DC, the dataset contained 1,853 cyclist-involved crashes and 2,344 bike lanes (21). Because the dataset was of moderate size, a brute-force nearest-neighbor search was used to compute the distance from each crash to every bike lane, requiring approximately 4.34 million point-to-polyline distance calculations. For each crash, the minimum perpendicular distance to a bike lane was retained.

A threshold of 100 m was used to distinguish crashes occurring near existing bicycle infrastructure from those occurring farther away. This distance was chosen to encompass crashes occurring on or reasonably close to bike lanes, where bicycle infrastructure was readily accessible. Crashes within 100 m of a bike lane were excluded from further analysis because nearby infrastructure already existed. Approximately 800 crashes remained after filtering. These crashes were then assigned to DDOT roadway centerlines to constrain clustering to individual roadway segments, preventing crashes on nearby but separate streets from being merged into the same cluster. Hierarchical complete-linkage clustering was subsequently applied using a 100 m linkage threshold, identifying localized groups of cyclist-involved crashes that may warrant further bicycle infrastructure review.

2.5 Normalization by exposure

Although the severity-sum metric identifies clusters with the greatest injury burden, it does not account for differences in traffic exposure. Roadways with higher traffic volumes are expected to experience more crashes simply because they carry more vehicles. To account for this effect, all clusters across the three crash categories were normalized by exposure using Annual Average Daily Traffic (AADT) values from DDOT's 2023 Traffic Volume dataset (22). According to Open Data DC, these values are derived from traffic counts collected on a three-year cycle and converted to AADT estimates.

Cluster locations and AADT roadway geometries were projected into a common coordinate reference system, and each cluster was assigned the nearest AADT location using Euclidean distance. Clusters located more than 100 m from the nearest AADT measurement were excluded because a representative traffic volume could not be assigned reliably. Five-year traffic exposure was estimated as AADT × 365 × 5 to approximate the total traffic volume over the study period. A traffic-normalized risk metric was then calculated by dividing the severity sum by the five-year exposure. This risk metric is multiplied by 100,000,000 to represent severity score per 100,000,000 vehicle exposures. For each cluster, a rank-change value was also calculated as the difference between its severity-sum rank and its exposure-normalized rank. Positive values indicate that a cluster moved higher in priority after exposure normalization, whereas negative values indicate that it moved lower. Finally, Spearman's rank correlation coefficient and the Top-10 Overlap metric were used to compare the exposure-normalized rankings with the original severity-sum rankings.

2.6 DBSCAN

To evaluate the robustness of CrashMap360's hotspot rankings, complete-linkage clustering was compared with DBSCAN, an unsupervised density-based clustering algorithm. DBSCAN groups nearby crashes based on two parameters: the minimum number of points required to form a cluster and the maximum distance between neighboring crashes (23). For this study, the minimum cluster size was set to two crashes, and the neighborhood distance matched the complete-linkage threshold for each crash category: 30 m for nighttime crashes, 59.4 m for speeding-related crashes, and 100 m for cyclist-involved crashes.

To compare hotspot locations across methods, complete-linkage and DBSCAN clusters were considered to represent the same location if their centroids were within 100 m of one another. This threshold was selected based on the work of Rahman et al., who used 100 m buffers to evaluate geographic factors associated with severe motor vehicle crashes (24).

Because DBSCAN and complete linkage use fundamentally different clustering strategies, they often produce different numbers of clusters and cluster boundaries. Consequently, one-to-one rank comparisons using Spearman's rank correlation coefficient are inappropriate. Instead, agreement between methods was evaluated using the Top-5, Top-10, and Top-20 Overlap metrics, which quantify the percentage of high-priority hotspots identified by both approaches. For nighttime crashes, Top-16 Overlap was used in place of Top-20 Overlap because only 16 complete-linkage clusters were identified.

2.7 DDOT corridors

The severity-sum and exposure-normalized rankings were also compared with the District Department of Transportation's (DDOT) High Injury Network (25). To determine whether a CrashMap360 cluster corresponded to a DDOT high-injury corridor, each cluster was matched to the nearest corridor using a 100 m threshold. This threshold was used to determine whether the two methods identified the same general hotspot location.

Because the DDOT High Injury Network and CrashMap360 produce different numbers of hotspots and do not provide directly comparable rankings, Spearman's rank correlation coefficient was not appropriate. Instead, agreement was evaluated using Top-*k* Overlap and the percentage of DDOT high-injury corridors represented within the top-*k* CrashMap360 rankings. The same *k* values used in the DBSCAN comparison were applied here (Top-5, Top-10, and Top-20; Top-16 for nighttime crashes).

2.8 Sensitivity and robustness of thresholds

The clustering thresholds were selected based on expert consultation and published literature. All three thresholds were reviewed with a Vision Zero executive to assess their practical relevance. For nighttime crashes, the 30 m threshold is consistent with published estimates of the effective illumination range of streetlights (26). For speeding-related crashes, previous studies have reported that speed cameras may influence driver behavior over distances of up to 200 m (27). However, because CrashMap360 analyzes crashes within a dense urban environment, a shorter threshold of 59.4 m, corresponding to the median DDOT roadway-segment length, was adopted to reflect the scale of individual city blocks. To assess the sensitivity of this assumption, an additional analysis was performed using a 200 m threshold. For cyclist-involved crashes, no quantitative guidance on the effective influence distance of bike lanes was identified in the literature; therefore, a 100 m threshold was selected based on the methodological considerations described previously.

Threshold sensitivity was evaluated for each crash category by perturbing the primary threshold by ±25%. Nighttime crash thresholds of 22.5 m and 37.5 m, speeding thresholds of 44.55 m and 74.25 m (in addition to 200 m), and cyclist thresholds of 75 m and 125 m were tested. Because changing the clustering threshold alters both the number and composition of clusters, clusters generated under different thresholds were matched using a 100 m centroid-distance criterion. Matched clusters were then compared using the percentage of matched clusters and Spearman's rank correlation coefficient after excluding unmatched clusters.

2.9 Predictors of high severity hotspots

The clustering and ranking framework constitutes the primary methodology of CrashMap360. However, these methods identify priority hotspots without indicating which crash characteristics are most strongly associated with severe outcomes. Therefore, logistic regression was performed as a supplementary interpretability analysis to examine associations between cluster characteristics and crash severity.

Using the infrastructure-filtered clusters from the primary analysis would bias these associations because each crash category excludes crashes occurring near its corresponding infrastructure. To avoid this bias, a separate clustering dataset was constructed. Crash data were filtered using the same procedures described previously, except that crashes near existing infrastructure were retained. Crashes were then clustered using hierarchical complete-linkage clustering with a 100 m threshold, which captures nearby crashes within the same general roadway area while maintaining localized clusters suitable for associational analysis. The 100 m threshold is supported by Rahman et al. (24). Each cluster was assigned four outcome measures: severity sum, average severity per crash, severe crash percentage, and exposure-normalized severity.

Nine predictors were evaluated across four logistic regression models. Six predictors described the composition and infrastructure characteristics of each cluster: percentage of nighttime crashes, percentage of speeding-related crashes, percentage of cyclist-involved crashes, median distance to the nearest streetlight, median distance to the nearest speed-control device, and median distance to the nearest bike lane. Three additional predictors represented interaction terms: the percentage of crashes involving both nighttime and speeding, nighttime and cyclist involvement, and speeding and cyclist involvement. Because crashes satisfying both conditions are relatively uncommon, these interaction terms were expected to exhibit weaker associations.

Separate logistic regression models were developed for four binary outcomes. A cluster was classified as positive if it ranked within the top 25% for (1) severity sum, (2) average severity per crash, (3) severe crash percentage, or (4) exposure-normalized severity. These models were used to identify cluster characteristics associated with the highest-severity hotspots rather than to predict individual crashes.

3. Results

3.1 Night accidents

The comparison between the severity-sum and capped-severity-sum rankings demonstrated strong agreement, with a Spearman rank correlation coefficient of 0.9382, a 100% Top-5 Overlap, and a 90% Top-10 Overlap. These results indicate that capping the contribution of individual crashes has little effect on the prioritization of nighttime hotspots, suggesting that the highest-ranked locations are driven primarily by repeated severe crashes rather than isolated extreme events. Comparisons between the severity-sum and crash-frequency rankings also showed substantial agreement ($\rho = 0.7735$), although the lower Top-5 Overlap (60%) indicates that ranking locations by injury severity identifies some different high-priority hotspots than ranking them by crash frequency alone. The Top-10 Overlap of 90% suggests that the two approaches identify many of the same candidate locations but prioritize them differently.

The highest-ranked nighttime hotspot had a severity sum of 12 across three crashes and was located on Magazine Road Southwest. The second-ranked hotspot had a severity sum of 11, also across three crashes, and was located on 22nd Street Northeast. The remaining highest-priority candidate locations for streetlight review are summarized in Table 2.

Comparing the severity-sum rankings with the exposure-normalized (risk) rankings yielded a Spearman rank correlation coefficient of 0.8118, with 100% Top-5 Overlap and 70% Top-10 Overlap. These results indicate that accounting for traffic exposure changes the ordering of some hotspots but largely preserves the highest-priority locations. Thus, the locations with the greatest overall injury burden also tend to remain important after normalizing for traffic volume.

Comparison with DBSCAN demonstrated moderate agreement between clustering methods. The Top-5, Top-10, and Top-16 Overlap values were 60%, 50%, and 37.5%, respectively, indicating that hotspot prioritization is sensitive to the clustering algorithm used. Although both methods identify areas of elevated crash concentration, complete-linkage clustering and DBSCAN produce noticeably different hotspot locations and rankings.

Agreement with the District Department of Transportation's (DDOT) High Injury Network was limited. For the severity-sum rankings, the Top-5, Top-10, and Top-16 Overlap values were 0%, 20%, and 18.8%, respectively. The corresponding overlaps for the exposure-normalized rankings were 0%, 10%, and 18.8%.

These low overlap values suggest that CrashMap360 identifies a substantially different set of candidate locations for streetlight review than DDOT's corridor-based High Injury Network, reflecting the different objectives and methodologies of the two approaches.

Table 2. The top 10 streetlight locations based on severity scoring.

	N CRASHES	SEVERITY SUM	CAPPED SEVERITY SUM	AVG SEVERITY	SEVERE CRASH PERCENTAGE	FATAL CRASH PERCENTAGE	AVG LON	AVG LAT
0	3	12.0	9	4.000000	66.666667	0.0	-77.017848	38.828147
1	3	11.0	6	3.666667	33.333333	0.0	-76.973359	38.889910
2	4	11.0	4	2.750000	0.000000	0.0	-76.963688	38.852248
3	3	7.0	6	2.333333	33.333333	0.0	-77.011737	38.906000
4	4	6.0	4	1.500000	0.000000	0.0	-77.011924	38.881656
5	3	5.0	3	1.666667	0.000000	0.0	-77.018431	38.820787
6	3	5.0	3	1.666667	0.000000	0.0	-76.957523	38.920742
7	4	5.0	4	1.250000	0.000000	0.0	-76.980346	38.900194
8	3	4.0	3	1.333333	0.000000	0.0	-76.959220	38.869424
9	3	4.0	3	1.333333	0.000000	0.0	-77.011502	38.881660

Table 3. The top 10 streetlight locations normalized by exposure.

	RANK CHANGE*	N CRASHES	SEVERITY SUM	RISK PER 100M**	AVG LON	AVG LAT
0	1	3	11.0	86.962881	-76.973359	38.889910
1	1	4	11.0	55.690633	-76.963688	38.852248
2	-2	3	12.0	55.106792	-77.017848	38.828147
3	0	3	7.0	35.964524	-77.011737	38.906000
4	0	4	6.0	35.515515	-77.011924	38.881656
5	0	3	5.0	29.766689	-77.018431	38.820787
6	3	3	4.0	23.677010	-77.011502	38.881660
7	5	3	3.0	13.776698	-77.018233	38.822512
8	0	3	4.0	12.440577	-76.959220	38.869424
9	2	4	4.0	11.929357	-76.993201	38.827968

*Positive values indicate that a cluster moved to a higher priority after exposure normalization; negative values indicate that it moved to a lower priority. **Severity score per 100,000,000 vehicle exposures.

3.2 Speeding

The comparison between the severity-sum and capped-severity-sum rankings showed very strong overall agreement, with a Spearman rank correlation coefficient of 0.9851. However, the Top-10 Overlap was only 40%, indicating that although the overall ordering of clusters was similar, the highest-priority locations changed substantially after capping the influence of individual crashes. This result suggests that several top-ranked clusters in the severity-sum ranking were driven by isolated crashes with exceptionally high injury severity, whereas the capped-severity metric placed greater emphasis on locations with repeated severe crashes. Comparisons between the severity-sum and crash-frequency rankings also demonstrated strong overall agreement ($\rho = 0.9596$), yet the Top-10 Overlap was only 20%, indicating that ranking hotspots by injury severity identifies substantially different priority locations than ranking them by crash frequency alone.

The highest-ranked speeding hotspot had a severity sum of 29 and consisted of a single crash on Connecticut Avenue. The second-ranked hotspot had a severity sum of 26, also consisting of a single crash, and was located on Massachusetts Avenue. In contrast, lower-ranked clusters within the top 10 contained multiple crashes with lower average severity, illustrating how the severity-sum metric can prioritize isolated catastrophic crashes over locations with recurring serious crashes. The remaining highest-priority candidate locations for speed-camera or speed-hump review are summarized in Tables 4 and 5.

Table 4. The top 10 speed bump/speed camera locations with the highest severity sums.

	N CRASHES	SEVERITY SUM	CAPPED SEVERITY SUM	AVG SEVERITY	SEVERE CRASH PERCENTAGE	FATAL CRASH PERCENTAGE	AVG LON	AVG LAT
0	1	29.0	7	29.00	100.0	100.0	-77.074797	38.963568
1	1	26.0	7	26.00	100.0	100.0	-77.034270	38.906647
2	15	24.0	15	1.60	0.0	0.0	-77.008611	38.892054
3	1	23.0	7	23.00	100.0	100.0	-77.055274	38.929997
4	4	23.0	16	5.75	75.0	25.0	-77.022157	38.882069
5	2	17.0	8	8.50	50.0	50.0	-76.960459	38.869896
6	2	16.0	8	8.00	50.0	50.0	-76.996185	38.879451
7	1	15.0	7	15.00	100.0	100.0	-76.949160	38.895133
8	1	15.0	7	15.00	100.0	100.0	-76.958788	38.930031
9	1	15.0	7	15.00	100.0	100.0	-77.009333	38.920590

Comparison between the severity-sum and exposure-normalized (risk) rankings yielded a Spearman rank correlation coefficient of 0.5096 and a Top-10 Overlap of only 10%. These findings indicate that accounting for traffic exposure substantially changes hotspot prioritization, suggesting that several locations with high overall injury burden experience correspondingly high traffic volumes, whereas other locations become higher priorities after normalization for exposure.

Comparison with DBSCAN demonstrated moderate agreement between clustering methods. The Top-5, Top-10, and Top-20 Overlap values were 60%, 50%, and 45%, respectively, indicating that hotspot identification is sensitive to the clustering

algorithm used. Although both approaches produce different hotspot boundaries and identify areas of elevated crash concentration, prioritizations.
complete-linkage clustering and DBSCAN

Table 5. The top 10 speed bump/speed camera locations with the highest capped severity sums.

	N CRASHES	SEVERITY SUM	CAPPED SEVERITY SUM	AVG SEVERITY	SEVERE CRASH PERCENTAGE	FATAL CRASH PERCENTAGE	AVG LON	AVG LAT
0	4	23.0	16	5.750000	75.000000	25.000000	-77.022157	38.882069
1	15	24.0	15	1.600000	0.000000	0.000000	-77.008611	38.892054
2	3	12.0	12	4.000000	66.666667	33.333333	-77.027256	38.917001
3	3	10.0	9	3.333333	33.333333	33.333333	-76.970316	38.897351
4	3	10.0	9	3.333333	66.666667	0.000000	-76.978137	38.849197
5	3	9.0	9	3.000000	33.333333	33.333333	-77.018098	38.814772
6	2	17.0	8	8.500000	50.000000	50.000000	-76.960459	38.869896
7	2	16.0	8	8.000000	50.000000	50.000000	-76.996185	38.879451
8	2	14.0	8	7.000000	50.000000	50.000000	-76.939001	38.869791
9	2	12.0	8	6.000000	50.000000	50.000000	-77.007322	38.836306

Table 6. The top 10 speed bump/speed camera locations normalized by exposure.

	RANK CHANGE*	N CRASHES	SEVERITY SUM	RISK PER 100M	AVG LON	AVG LAT
0	24	1	10.0	626.223092	-77.008624	38.970264
1	12	1	13.0	378.093825	-76.985509	38.893392
2	4	2	16.0	331.962260	-76.996185	38.879451
3	47	1	8.0	306.542772	-76.988277	38.905470
4	72	1	7.0	252.675655	-76.975921	38.928424
5	34	2	9.0	214.041096	-77.011928	38.944409
6	102	1	5.0	191.589233	-76.987750	38.906592
7	168	1	4.0	183.720102	-76.973184	38.869476
8	47	1	8.0	169.314857	-76.927915	38.887207
9	19	3	10.0	169.171104	-76.970316	38.897351

*A greater magnitude of rank changes is reflected in a Spearman Rank Correlation Coefficient of 0.5096 which shows a moderate difference between the results from the two rankings methods. These differences exist because of the influence of traffic volume in the severity ranked clusters. These differences highlight the need for both top 10 sets of clusters to be examined. There are no negative values because cluster locations that were previously in the top 10 of severity ranked clusters do not appear in the top 10 of exposure normalized crashes.

Agreement between CrashMap360 and the DDOT High Injury Network depended on the ranking metric used. The severity-sum rankings exhibited Top-5, Top-10, and Top-20 Overlap values of 80%, 80%, and 65%, respectively, indicating substantial agreement with DDOT's corridor-based prioritization. In contrast, the exposure-normalized rankings produced overlap values of 20% for all three comparisons, demonstrating that normalizing by traffic volume identifies a markedly different set of candidate locations for speed-related infrastructure review.

3.3 Bike lanes

The comparison between the severity-sum and capped-severity-sum rankings demonstrated extremely strong agreement, with a Spearman rank correlation coefficient of 0.9891 and 100% Top-5 and Top-10 Overlap. These findings indicate that capping the contribution of individual crashes has virtually no effect on hotspot prioritization for cyclist-involved crashes. Unlike the speeding analysis, the highest-ranked cyclist hotspots are driven primarily by repeated severe crashes rather than isolated catastrophic events. Comparisons between the severity-sum and crash-frequency rankings also showed substantial agreement ($\rho = 0.8557$), although the Top-5 and Top-10 Overlap values of 60% indicate that injury severity provides additional information beyond crash frequency when prioritizing locations for bicycle infrastructure review.

The highest-ranked cyclist hotspot had a severity sum of 17 across eight crashes and was located on N Street Northwest. The second-ranked hotspot had a severity sum of 13 across six crashes near Scott Circle. In contrast to the speeding analysis, the highest-ranked cyclist hotspots consisted of multiple crashes rather than single severe events, suggesting that these locations represent persistent safety concerns. The remaining highest-priority candidate locations for bicycle infrastructure review are summarized in Table 7.

Comparison between the severity-sum and exposure-normalized (risk) rankings yielded a Spearman rank correlation coefficient of 0.5200, with Top-5 and Top-10 Overlap values of 40% and 50%, respectively. These results indicate that accounting for traffic exposure substantially alters hotspot prioritization, suggesting that several cyclist hotspots assume higher or lower priority once differences in roadway usage are considered.

Comparison with DBSCAN demonstrated moderate agreement between clustering methods. The Top-5, Top-10, and Top-20 Overlap values were 40%, 70%, and 60%, respectively, indicating that hotspot identification is influenced by the clustering algorithm but that many of the highest-priority cyclist hotspots are identified by both methods.

Agreement between CrashMap360 and the DDOT High Injury Network was limited. The severity-sum rankings produced Top-5, Top-10, and Top-20 Overlap values of 20%, 20%, and 35%, respectively, while the exposure-normalized rankings yielded 20%, 10%, and 30% overlap. These relatively low agreement values indicate that CrashMap360 identifies a substantially different set of candidate locations for bicycle infrastructure review than DDOT's corridor-based High Injury Network, regardless of whether rankings are based on severity or exposure-normalized risk.

Table 7. The top 10 locations for bike lanes ranked by severity scoring.

	N CRASHES	SEVERITY SUM	CAPPED SEVERITY SUM	AVG SEVERITY	SEVERE CRASH PERCENTAGE	FATAL CRASH PERCENTAGE	AVG LON	AVG LAT
0	8	17.0	17	2.125000	37.500000	0.0	-77.022012	38.920637
1	6	13.0	12	2.166667	33.333333	0.0	-77.036632	38.907261
2	3	9.0	9	3.000000	66.666667	0.0	-77.001728	38.900205
3	3	9.0	9	3.000000	66.666667	0.0	-77.059818	38.936757
4	5	8.0	8	1.600000	20.000000	0.0	-77.024321	38.936515
5	5	8.0	8	1.600000	20.000000	0.0	-77.009379	38.909669
6	4	7.0	7	1.750000	25.000000	0.0	-77.061334	38.905214
7	4	7.0	7	1.750000	25.000000	0.0	-77.029499	38.917048
8	3	6.0	6	2.000000	33.333333	0.0	-77.005888	38.907687
9	3	6.0	6	2.000000	33.333333	0.0	-77.039365	38.901605

Table 8. The top 10 locations for bike lanes normalized by exposure.

	RANK CHANGE	N CRASHES	SEVERITY SUM	RISK PER 100M	AVG LON	AVG LAT
0	3	3	9.0	82.150705	-77.059818	38.936757
1	18	3	4.0	75.267199	-77.026112	38.935023
2	11	3	6.0	74.804806	-77.037914	38.930150
3	2	5	8.0	52.858575	-77.009379	38.909669
4	-3	6	13.0	52.384819	-77.036632	38.907261
5	20	4	4.0	51.534936	-77.013752	38.927761
6	2	3	6.0	49.798110	-77.005888	38.907687
7	34	3	3.0	44.236696	-77.008508	38.905644
8	-8	8	17.0	38.107791	-77.022012	38.920637
9	1	3	6.0	31.679237	-77.049339	38.909415

3.4 Maps

In Figure 1, the clusters with the highest severity sums are plotted with darker and larger circles corresponding to clusters with higher severity sums. As such, the larger, orange cluster in Figure 1 has a higher severity score than the two tan circles. All the cyclist accidents far from bike lanes are also plotted in blue on the map. The same structure is used across the other two categories, with the main difference only being a different crash type for each map.

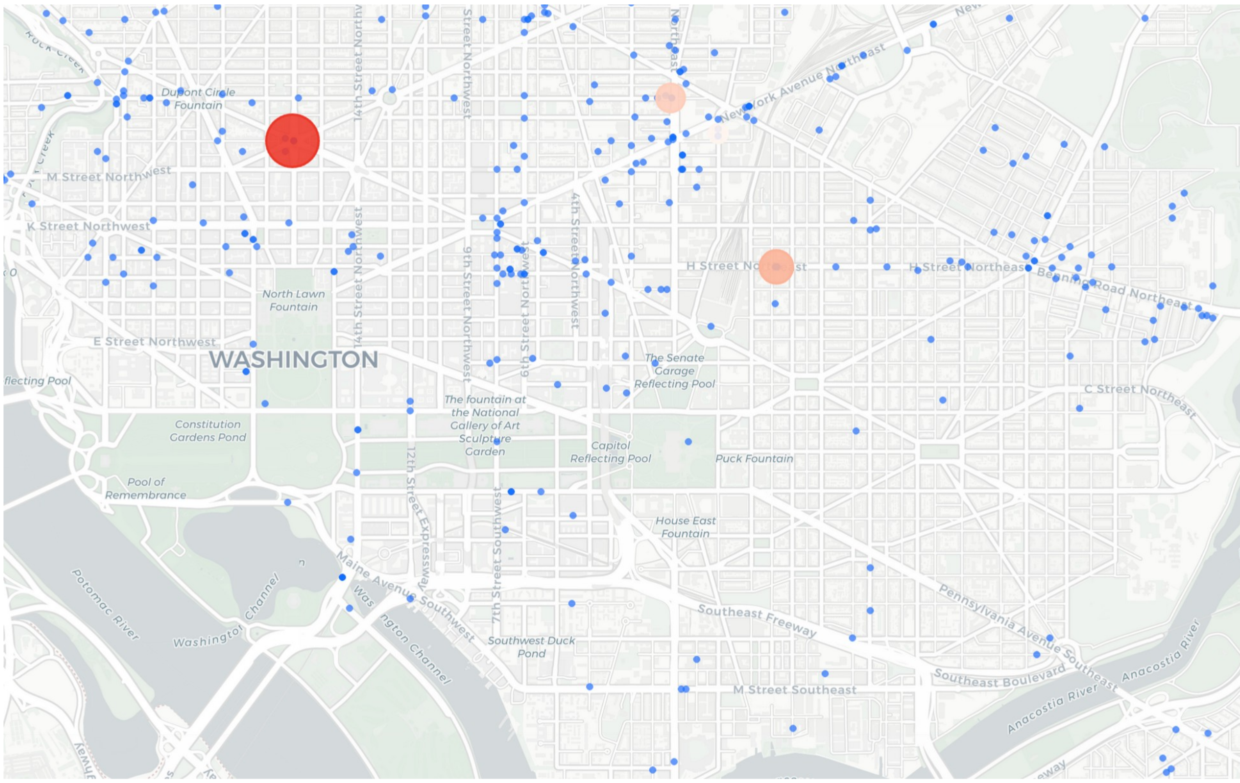

Figure 1. A map of all the candidate locations for bike-lane review.

3.5 Sensitivity and robustness

Threshold sensitivity was evaluated by comparing the primary clustering threshold for each crash category with alternative threshold values. Overall, the rankings remained highly consistent across all threshold perturbations, indicating that CrashMap360 is robust to reasonable changes in clustering distance.

For nighttime crashes, the primary 30 m threshold demonstrated strong robustness. Reducing the threshold to 22.5 m produced 50.0% matched clusters and a Spearman rank correlation coefficient of 0.9794, while increasing the threshold to 37.5 m increased the matched clusters to 81.25% with a Spearman rank correlation coefficient of 0.9271. Although changes in the clustering threshold altered the composition of individual clusters, the high rank correlations indicate that the relative prioritization of hotspot locations remained largely unchanged.

For speeding-related crashes, the primary 59.4 m threshold was the most robust of the three crash-category analyses. Thresholds of 44.55 m and 74.25 m produced 99.8% and 99.1% matched clusters, respectively, with corresponding Spearman rank correlation coefficients of 0.9973 and 0.9935. Even when the clustering threshold was increased substantially to 200 m, consistent with previous literature, agreement remained high (92.5% matched clusters; ρ = 0.9543). These findings indicate that hotspot prioritization for speeding-related crashes is largely insensitive to reasonable changes in clustering distance.

For cyclist-involved crashes, the primary 100 m threshold also demonstrated strong robustness. Reducing the threshold to 75 m produced 93.5% matched clusters and a Spearman rank correlation coefficient of 0.8974, whereas increasing the threshold to 125 m resulted in 90.0% matched clusters and a Spearman rank correlation coefficient of 0.7182. Although increasing the threshold altered the ordering of some hotspots more substantially than in the other crash categories, the high matched cluster percentages indicate that most hotspot locations remained stable.

Overall, the sensitivity analyses demonstrate that CrashMap360 produces stable hotspot locations and generally consistent hotspot rankings across a range of plausible clustering thresholds. The robustness was strong for speeding-related crashes and remained substantial for nighttime and cyclist-involved crashes, supporting the use of the selected thresholds in the primary analyses.

3.6 Predictors of high severity hotspots

The logistic regression analyses examined which crash and infrastructure characteristics were associated with clusters in the highest quartile of severity sum, average severity per crash, severe crash percentage, and exposure-normalized severity. The results, including odds ratios and *p* values for all predictors, are summarized in Table 9.

Within the severity-sum model, the strongest positive associations were observed for the percentage of nighttime crashes (OR = 1.21), the percentage of speeding-related crashes (OR = 1.22), and median distance to the nearest bike lane (OR = 1.11), with all three predictors achieving statistical significance ($p < 0.05$). In contrast, a higher percentage of cyclist-involved crashes (OR = 0.78) and greater distance to the nearest speed-control device (OR = 0.76) were associated with lower odds of belonging to the highest-severity clusters. These findings suggest that nighttime crashes, speeding-related crashes, and locations farther from bicycle infrastructure are more likely to occur within the most severe hotspots.

Within the average severity per crash model, the percentage of speeding-related crashes exhibited the strongest positive association (OR = 1.34), followed by the percentage of nighttime crashes (OR = 1.10) and distance to the nearest bike lane (OR = 1.21). These results indicate that clusters containing more speeding-related crashes and those farther from bicycle infrastructure tend to have higher average crash severity.

The severe crash percentage and exposure-normalized severity models similarly identified speeding-related crashes as the strongest positive predictor, indicating that speeding was consistently associated with more severe crash outcomes regardless of the severity metric considered.

Across all four models, the three interaction predictors—nighttime and speeding, nighttime and cyclist involvement, and speeding and cyclist involvement—generally exhibited *p* values greater than 0.05, providing limited statistical evidence that these combined crash characteristics contributed additional explanatory value beyond the individual predictors.

Table 9. The odds ratio and p values of the logistic regression predictors.

Predictor	Severity Sum	Average Severity	Serious Crash Percentage	Exposure-Normalized Severity
Night crash percentage	1.21, $p < 0.001$	1.10, $p = 0.060$	1.13, $p = 0.020$	1.07, $p = 0.221$
Speeding crash percentage	1.22, $p = 0.003$	1.34, $p < 0.001$	1.19, $p = 0.009$	1.15, $p = 0.044$
Bike-involved crash percentage	0.78, $p = 0.011$	0.85, $p = 0.069$	1.01, $p = 0.854$	0.91, $p = 0.254$
Night and speeding crash percentage	0.92, $p = 0.236$	0.96, $p = 0.574$	1.01, $p = 0.899$	0.88, $p = 0.080$
Night and bike-involved crash percentage	0.95, $p = 0.585$	0.95, $p = 0.548$	1.01, $p = 0.844$	0.94, $p = 0.443$
Speeding and bike-involved crash percentage	0.95, $p = 0.431$	0.94, $p = 0.304$	0.96, $p = 0.441$	0.94, $p = 0.343$
Median distance to nearest streetlight	1.06, $p = 0.227$	1.09, $p = 0.100$	1.10, $p = 0.046$	0.93, $p = 0.245$
Median distance to nearest speed-related device	0.76, $p < 0.001$	1.04, $p = 0.403$	1.06, $p = 0.242$	1.06, $p = 0.230$
Median distance to nearest bike lane	1.11, $p = 0.031$	1.21, $p < 0.001$	1.05, $p = 0.350$	0.97, $p = 0.534$

4. Discussion

CrashMap360 presents a multi-stage framework for prioritizing candidate roadway safety interventions by integrating infrastructure-proximity screening, severity-based hotspot ranking, exposure normalization, robustness testing, and supplementary interpretability analyses. Unlike conventional hotspot analyses that primarily identify locations with large numbers of crashes, CrashMap360 prioritizes locations where severe crash patterns occur beyond the expected influence of existing infrastructure. Rather than prescribing specific engineering interventions, the framework is intended to reduce the search space for transportation engineers by identifying a manageable set of candidate locations that warrant detailed field investigation.

Several broad findings emerged across the three crash categories. First, severity-based rankings consistently provided information beyond crash frequency by emphasizing locations with greater injury burden rather than simply larger numbers of crashes. Second, the influence of exposure normalization differed by crash category, producing relatively small changes in nighttime crash prioritization but substantially altering the rankings of speeding- and cyclist-related hotspots.

Third, robustness analyses demonstrated that the framework was generally insensitive to reasonable changes in clustering thresholds, indicating that the recommended locations were driven primarily by the underlying spatial distribution of crashes rather than specific parameter choices. Finally, the comparison with DDOT's High Injury Network showed that CrashMap360 identifies a different but complementary set of candidate locations by focusing on localized infrastructure screening rather than corridor-level crash burden. Together, these findings suggest that

CrashMap360 functions as a practical decision-support framework that complements existing transportation safety analyses by identifying candidate locations for further engineering evaluation.

4.1 Streetlights

The nighttime crash analysis demonstrates that hotspot prioritization is largely insensitive to the choice of ranking metric. The high Spearman rank correlation coefficient ($\rho = 0.9382$) together with the high Top-5 (100%) and Top-10 (90%) Overlap between the severity-sum and capped-severity-sum rankings indicate that limiting the influence of individual crashes has little effect on hotspot prioritization. Unlike the speeding analysis, the highest-ranked nighttime hotspots are driven primarily by repeated severe crashes rather than isolated catastrophic events. This interpretation is further supported by Table 2, where every top-10 hotspot contains at least three crashes and none contain fatalities, indicating that persistent crash patterns rather than singular extreme crashes dominate the rankings.

The comparison between severity and crash frequency ($\rho = 0.7735$) indicates that the two measures are closely related but not interchangeable. Although many high-priority locations appear in both rankings, the severity-based approach captures information beyond crash counts by giving greater weight to locations where crashes result in more serious injuries. Consequently, severity ranking provides a more informative basis for prioritizing candidate streetlight review locations than crash frequency alone.

The relatively high agreement between the severity-sum and exposure-normalized rankings demonstrates that accounting for traffic volume produces only modest changes in hotspot prioritization. The complete retention of the Top-5 hotspots and the relatively high Spearman correlation suggest that the highest-priority nighttime locations remain important even after differences in roadway exposure are considered. This consistency indicates that the identified hotspots reflect genuinely elevated nighttime crash burden rather than simply high traffic volume.

The moderate agreement between complete-linkage clustering and DBSCAN indicates that hotspot identification is influenced by the clustering algorithm. However, complete-linkage clustering is better suited to CrashMap360 because it produces compact, localized clusters that correspond more closely to candidate engineering intervention areas and avoids chaining together crashes over larger distances.

The limited agreement between CrashMap360 and the DDOT High Injury Network indicates that the streetlight-screening approach identifies a different set of candidate locations than DDOT's corridor-based framework. This difference likely reflects the distinct objectives of the two methods: DDOT prioritizes corridors with historically high crash burdens, whereas CrashMap360 specifically identifies clusters of nighttime crashes occurring beyond the screening distance from existing streetlights. Consequently, the two approaches should be viewed as complementary rather than competing prioritization strategies.

Finally, the streetlight analysis evaluates proximity to existing streetlights rather than lighting quality. CrashMap360 does not measure illumination levels, tree canopy obstruction, fixture performance, roadway geometry, or other environmental factors that influence nighttime visibility. Therefore, identified hotspots should be considered candidate locations for field investigation rather than direct recommendations for installing additional streetlights.

4.2 Speed cameras and speed bumps

The speeding analysis demonstrates that hotspot prioritization depends strongly on the metric used. Unlike the nighttime and cyclist analyses, the highest-ranked speeding hotspots change substantially when alternative prioritization methods are applied, indicating that different engineering objectives lead to different candidate intervention locations.

The severity-sum rankings are heavily influenced by isolated catastrophic crashes. As shown in Table 4, many of the highest-ranked clusters consist of only a single crash, indicating that one exceptionally severe collision can outweigh locations with repeated serious crashes. This effect is reflected in the relatively low Top-10 Overlap (40%) between the severity-sum and capped-severity-sum rankings. However, the high Spearman rank correlation coefficient (>0.95) indicates that the overall ordering of hotspots remains similar despite changes among the highest-ranked locations. The capped-severity metric therefore provides a useful complementary perspective by reducing the influence of isolated catastrophic crashes and emphasizing locations with recurring severe crash patterns that may represent persistent roadway safety deficiencies. The comparison between crash frequency and severity rankings further demonstrates that the most frequently occurring speeding crashes are not necessarily the most severe. Only 2 of the top 10 hotspots were shared between the two rankings, although the high Spearman rank correlation coefficient indicates that crash frequency and severity remain closely related across the complete set of clusters. These findings suggest that relying exclusively on crash counts may overlook locations where severe injuries occur disproportionately often.

Exposure normalization produced the greatest change in hotspot prioritization. The relatively low Spearman rank correlation coefficient ($\rho = 0.5096$) and Top-10 Overlap (10%) indicate that accounting for traffic volume substantially alters which locations emerge as the highest priorities. Consequently, exposure-normalized rankings emphasize crash risk relative to roadway usage rather than total crash burden, potentially elevating locations with fewer crashes but disproportionately high injury rates while reducing the priority of heavily traveled corridors where large crash totals primarily reflect greater traffic exposure.

The moderate agreement between complete-linkage clustering and DBSCAN indicates that hotspot identification is influenced by the clustering algorithm, although many high-priority locations are identified by both methods. Comparison with the DDOT High Injury Network provides a similar perspective. The relatively high overlap between the severity-based rankings and DDOT's corridor prioritization suggests that both approaches identify many of the same high-burden locations. In contrast, the much lower agreement between the exposure-normalized rankings and the DDOT network indicates that incorporating traffic volume produces a substantially different prioritization strategy focused on traffic-adjusted crash risk rather than absolute crash burden.

The identified hotspots should therefore be interpreted as candidate locations for engineering review rather than direct recommendations for speed cameras or speed humps. Infrastructure proximity serves as a screening criterion rather than evidence of causation, and selecting an appropriate intervention requires consideration of roadway geometry, operating speeds, traffic volumes, enforcement history, and other site-specific factors that are not captured by the model.

4.3 Bike lanes

The extremely high Spearman rank correlation coefficient and the 100% Top-5 and Top-10 Overlap between the uncapped and capped severity rankings indicate that limiting the contribution of individual crashes has virtually no effect on cyclist hotspot prioritization. Unlike the speeding analysis, cyclist hotspots are generally driven by multiple recurring crashes rather than isolated catastrophic events. This interpretation is further supported by Table 7, where none of the top 10 cyclist clusters contained fatalities. These findings suggest that cyclist hotspot identification should not rely solely on fatal crashes, as locations with repeated moderate-severity injuries may represent persistent infrastructure deficiencies that warrant engineering attention.

The relatively high Spearman rank correlation coefficient ($\rho = 0.8557$) between the frequency-based and severity-based rankings indicates that crash frequency and injury severity are closely related for cyclist crashes. Consequently, agencies relying primarily on crash counts are unlikely to overlook many of the highest-severity cyclist hotspots, although severity-based ranking still provides additional information for prioritization.

The substantially lower Spearman rank correlation ($\rho = 0.5200$) between the severity and exposure-normalized rankings demonstrates that accounting for traffic exposure meaningfully changes hotspot prioritization. This finding suggests that locations with fewer crashes may represent greater relative risk once roadway usage is considered, emphasizing the importance of exposure normalization when prioritizing cyclist safety improvements.

The moderate agreement between complete-linkage clustering and DBSCAN indicates that hotspot identification remains somewhat dependent on the clustering algorithm, consistent with the findings for nighttime and speeding crashes. Similarly, the relatively low agreement between CrashMap360 and DDOT's High Injury Network suggests that the infrastructure-screening approach identifies a different set of candidate locations than DDOT's corridor-based prioritization. Rather than representing competing methods, these approaches provide complementary perspectives on cyclist safety, with CrashMap360 emphasizing localized infrastructure review.

Field investigations should therefore evaluate roadway geometry, cyclist volumes, intersection design, parking activity, turning movements, and existing traffic-control devices to determine whether adding a bike lane—or another engineering intervention—would provide the greatest safety benefit.

4.4 Sensitivity and robustness

The sensitivity analyses indicate that CrashMap360 is generally robust to reasonable changes in the infrastructure-distance and clustering thresholds. Across all three crash categories, modest threshold perturbations produced similar hotspot rankings, suggesting that the identified candidate locations are driven primarily by the underlying spatial distribution of crashes rather than by a particular parameter choice.

For nighttime crashes, reducing the primary threshold from 30 m to 22.5 m produced a relatively low cluster-match percentage (50%). This decrease is expected because lowering the threshold classifies many additional crashes as being far from existing streetlights, increasing the number of retained crashes and generating new clusters. Despite these structural changes, the Spearman rank correlation between matched clusters remained extremely high, indicating that the rankings of hotspot locations changed very little. Increasing the threshold to 37.5 m produced both a high cluster-match percentage and a high Spearman correlation, demonstrating that modest increases in the threshold have little effect on the resulting prioritization.

For speeding-related crashes, the results were particularly robust. Comparisons between the primary 59.4 m threshold and the 44.55 m and 74.25 m thresholds produced cluster-match percentages greater than 99% and Spearman rank correlation coefficients greater than 0.99, indicating that small changes in the clustering distance have a negligible effect on hotspot identification. Even when the threshold was increased to 200 m, consistent with previous literature, hotspot rankings remained highly similar. This robustness likely reflects the assignment of crashes to individual roadway segments, which constrains clustering to the relevant street and reduces sensitivity to moderate changes in clustering distance.

The cyclist analysis also demonstrated substantial robustness. Comparisons between the primary 100 m threshold and the 75 m threshold produced both high cluster-match percentages and high Spearman correlations, indicating that hotspots and their rankings remained largely unchanged. Increasing the threshold to 125 m maintained a high cluster-match percentage but reduced the Spearman correlation, suggesting that while most hotspot locations were preserved, their relative rankings changed. This reduction is likely due to additional crashes being incorporated into existing clusters as the threshold increased, altering cluster severity sums without substantially changing hotspot locations. The 100 m threshold therefore remains a reasonable balance between capturing crashes influenced by nearby bicycle infrastructure and maintaining localized, interpretable hotspot boundaries.

Overall, the sensitivity analyses demonstrate that CrashMap360 is not overly dependent on a single set of threshold values. Although modifying thresholds changes the composition of some clusters, prioritization remains largely consistent across analyses. These findings increase confidence that the framework identifies stable candidate locations for engineering review rather than artifacts created by specific parameter selections.

4.5 Predictors of high severity hotspots

After CrashMap360 identifies candidate locations through infrastructure-proximity screening and hotspot ranking, the logistic regression provides an interpretability analysis by identifying the crash characteristics most consistently associated with severe hotspots. Unlike the clustering and ranking framework, the logistic regression does not influence hotspot selection; instead, it helps explain why certain locations emerge as high priorities for engineering review.

Across all four logistic regression models, speeding-related crashes were the most consistent positive predictor of upper-quartile hotspot severity. This finding suggests that speeding is more strongly associated with severe crash clusters than the other characteristics examined and supports prioritizing speeding-related conditions during engineering review of CrashMap360 hotspots. Notably, speeding remained positively associated regardless of whether hotspot severity was defined by total severity, average severity per crash, severe crash percentage, or exposure-normalized severity, indicating that its association is robust across multiple definitions of crash severity.

Nighttime crashes and greater distance from bicycle infrastructure were also positively associated with severe hotspots, although these relationships were less consistent across the four models. These findings suggest that locations with frequent nighttime crashes or limited nearby bicycle infrastructure may warrant additional investigation, but the strength of these associations depends on the severity metric being considered.

The interaction predictors (nighttime and speeding, nighttime and cyclist involvement, and speeding and cyclist involvement) generally exhibited limited statistical support, indicating that combinations of crash characteristics provided little additional explanatory value beyond the individual predictors themselves.

These findings should be interpreted as associational rather than causal. The logistic regression identifies characteristics that commonly occur in severe hotspots but does not demonstrate that modifying a single infrastructure feature would reduce crash severity. Instead, the regression complements CrashMap360 by providing context for interpreting hotspot rankings and helping transportation engineers prioritize field investigations of the identified candidate locations.

4.6 Operational use by transportation agencies

Agencies such as the DDOT could use Crashmap360 as an initial screening tool. Top locations identified for review by Crashmap360 could be compared to high ranked corridors in the DDOT's High Injury Network. Locations from Crashmap360 not identified by the DDOT could be flagged for further review, while locations already identified could be used as reinforcement. Agency staff should conduct field review to understand traffic volumes, existing infrastructure, crash history, and land use to determine whether infrastructure intervention, greater enforcement, roadway design changes or no intervention would be most appropriate.

4.7 Limitations

The crash records obtained from Open Data DC contain limited information on the circumstances surrounding each crash because of privacy and reporting constraints. Consequently, CrashMap360 cannot account for many crash-specific factors, such as driver impairment, distraction, weather, vehicle condition, or detailed roadway characteristics. Incorporating additional crash-level variables in future work could improve interpretation of the factors contributing to hotspot formation.

CrashMap360 identifies associations rather than causal relationships. A hotspot located far from a particular type of infrastructure does not necessarily imply that the absence of that infrastructure caused the crashes or that installing it would reduce crash frequency or severity. Likewise, infrastructure placement is not random. Existing speed cameras, speed humps, bike lanes, and streetlights are often installed at locations with historically elevated crash risk. Consequently, differences between crashes occurring near and far from infrastructure may partly reflect underlying roadway characteristics, traffic patterns, land use, or previous engineering decisions rather than the effect of the infrastructure itself.

The accuracy of hotspot identification also depends on the quality of the underlying data. Crash reporting may be incomplete or inconsistent across locations, crash coordinates may contain positional errors, and some roadway or infrastructure datasets may be incomplete or outdated. These limitations can affect distance calculations, cluster formation, and hotspot rankings.

Several methodological choices require subjective judgment, including the selection of infrastructure-distance thresholds, clustering distances, and severity-weighting values. Although sensitivity analyses demonstrated that the primary findings were generally robust to reasonable threshold perturbations, no single threshold or weighting scheme should be considered universally optimal.

Exposure normalization is based on Annual Average Daily Traffic (AADT) estimates rather than direct measurements of traffic throughout the study period. AADT values from the 2023 DDOT Traffic Volume dataset were applied as approximations over the five-year analysis period and represent only motor vehicle exposure. They do not account for cyclist volume, pedestrian activity, temporal changes in traffic demand, or localized traffic conditions that may differ from the nearest AADT measurement location.

Recommended intervention locations should therefore be considered screening candidates rather than definitive engineering recommendations. Field inspections remain necessary because existing infrastructure may be absent from the available datasets, infrastructure effectiveness may be influenced by local environmental conditions such as vegetation or buildings, and roadway design characteristics not represented in the datasets may affect crash risk.

Finally, CrashMap360 has not yet undergone external validation against future crash outcomes or engineering interventions. Future research should evaluate whether the framework identifies hotspots more effectively than existing prioritization methods and whether its recommendations lead to measurable improvements in traffic safety following implementation.

5. Conclusion

Crashmap360 is an infrastructure-prioritization framework that analyzes Washington, D.C. crash data from the last five years to identify candidate locations for engineering review. The framework separates crashes into three categories—nighttime, speeding-involved, and cyclist crashes—and, within each category, calculates the distance between crashes and related protective infrastructure before identifying clusters of crashes occurring beyond predefined infrastructure-distance thresholds. These candidate locations are ranked using crash severity and exposure-normalized risk and are further evaluated against alternative prioritization methods, including capped-severity rankings, crash-frequency rankings, DBSCAN clustering, and the Department of Transportation's High Injury Network.

The results demonstrate that hotspot prioritization can vary substantially depending on the ranking methodology used, while sensitivity analyses show that the framework remains reasonably robust to changes in clustering thresholds. Rather than predicting future crashes or establishing causal relationships between infrastructure and crash occurrence, Crashmap360 provides a systematic, data-driven approach for identifying locations that may warrant further engineering investigation. The framework is intended to complement, rather than replace, existing transportation safety practices by providing an additional tool for infrastructure prioritization.

Although the study is limited to Washington, D.C. and has not yet been externally validated against future intervention outcomes, the methodology is generalizable to other municipalities with comparable crash, infrastructure, and traffic-volume datasets. Future work should evaluate the framework using independent transportation datasets and assess whether its prioritization improves engineering decision-making and safety outcomes relative to existing approaches.

Code and data repository

https://github.com/SachinKundra09/Crashmap360.git

The assistance of large language models was used in developing parts of Crashmap360's software. All design, algorithms, and methodology decisions, along with interpretations, were made by the author.

6. References

1. National Highway Traffic Safety Administration. (2025, April 8). *NHTSA estimates 39,345 traffic fatalities in 2024*. U.S. Department of Transportation. https://www.nhtsa.gov/press-releases/nhtsa-estimates-39345-traffic-fatalities-2024

2. Centers for Disease Control and Prevention. (2025, January 16). *Risk factors for teen drivers*. U.S. Department of Health and Human Services. https://www.cdc.gov/teen-drivers/risk-factors/index.html

3. National Highway Traffic Safety Administration. (2025, April). *Traffic safety facts 2023 data: Children* (Report No. DOT HS 813 712). U.S. Department of Transportation. https://crashstats.nhtsa.dot.gov/Api/Public/Publication/813712

4. Stein, P., & Hedgpeth, D. (2024, August 23). *Traffic fatalities have doubled since D.C.'s promise of zero*. *The Washington Post*. https://www.washingtonpost.com/dc-md-va/2024/08/23/vision-zero-dc-traffic-deaths/

5. The Lab @ DC. (2025). *Can targeting messages to high-risk drivers prevent future tickets and crashes?* Government of the District of Columbia. https://thelabprojects.dc.gov/high-risk-drivers

6. Weiner, R. (2025, September 3). *D.C. can predict who will get into car crashes but can't stop them*. *The Washington Post*. https://www.washingtonpost.com/dc-md-va/2025/09/03/dc-car-crash-text-experiment/

7. Strauss, T., & Lentz, J. (2009, March). *Spatial scale of clustering of motor vehicle crash types and appropriate countermeasures* (Final Report, MTC Project 2007-10; Iowa DOT Project TS-UNI 0701). Institute for Transportation, Iowa State University. https://www.intrans.iastate.edu/wp-content/uploads/2018/03/Crash_patterns.pdf

8. Hovenden, E., & Liu, G.-J. (2020). *Use of spatial analysis techniques to identify statistically significant crash hot spots in metropolitan Melbourne*. *Journal of Road Safety, 31*(4), 36–58. https://doi.org/10.33492/JRS-D-19-00249

9. Alam, M. S., & Tabassum, N. J. (2023). *Spatial pattern identification and crash severity analysis of road traffic crash hot spots in Ohio*. *Heliyon, 9*(5), e16303. https://doi.org/10.1016/j.heliyon.2023.e16303

10. Mansfield, T. J., Peck, D., Morgan, D., McCann, B., & Teicher, P. (2018). *The effects of roadway and built environment characteristics on pedestrian fatality risk: A national assessment*

at the neighborhood scale. *Accident Analysis & Prevention, 121*, 166–176. https://doi.org/10.1016/j.aap.2018.06.018

11. Staton, C., Vissoci, J. R. N., Gong, E., Toomey, N., Wafula, R., Abdelgadir, J., Zhou, Y., Liu, C., Pei, F., Zick, B., Ratliff, C. D., Rotich, C., Jadue, N., de Andrade, L., von Isenburg, M., Hocker, M., & Charles, A. (2016). *Road traffic injury prevention initiatives: A systematic review and metasummary of effectiveness in low- and middle-income countries*. *PLOS ONE, 11*(1), e0144971. https://doi.org/10.1371/journal.pone.0144971

12. Wang, D., Krebs, E., Vissoci, J. R. N., de Andrade, L., Rulisa, S., & Staton, C. A. (2020). *Built environment analysis for road traffic crash hotspots in Kigali, Rwanda*. *Frontiers in Sustainable Cities, 2*, Article 17. https://doi.org/10.3389/frsc.2020.00017

13. District of Columbia Government. (2025). *Crashes in DC* [Data set]. Open Data DC. https://opendata.dc.gov/datasets/70392a096a8e431381f1f692aaa06afd_24

14. Sharafeldin, M., Ksaibati, K., & Gerow, K. (2023). *Frequency analysis of equivalent property-damage-only (EPDO) crashes at intersections*. *Eng, 4*(2), 1116–1126. https://doi.org/10.3390/eng4020064

15. Spearman, C. (1904). *The proof and measurement of association between two things*. *The American Journal of Psychology, 15*(1), 72–101. https://doi.org/10.2307/1412159

16. District Department of Transportation. (2025). *Street lights* [Data set]. Open Data DC. https://opendata.dc.gov/datasets/street-lights

17. CU Phosco. (2025). *Why are streetlights so tall? Street Lighting 101*. https://www.cuphosco.com/articles/why-are-streetlights-so-tall

18. District Department of Transportation. (2023, November). *Automated traffic enforcement camera locations*. Government of the District of Columbia. https://ddot.dc.gov/publication/automated-traffic-enforcement-camera-locations

19. District Department of Transportation. (2023). *Speed humps* [Data set]. Open Data DC. https://opendata.dc.gov/datasets/DCGIS::speed-humps

20. District Department of Transportation. (2025). *Roadway centerlines* [Data set]. Open Data DC. https://opendata.dc.gov/pages/roadway-centerlines

21. District Department of Transportation. (2023). *Bicycle lanes* [Data set]. Open Data DC. https://opendata.dc.gov/datasets/bicycle-lanes

22. District Department of Transportation. (2024). *2023 traffic volume* [Data set]. Open Data DC. https://opendata.dc.gov/datasets/DCGIS::2023-traffic-volume/about

23. Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. (1996). *A density-based algorithm for discovering clusters in large spatial databases with noise*. In E. Simoudis, J. Han, & U. Fayyad (Eds.), *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96)* (pp. 226–231). AAAI Press. https://dl.acm.org/doi/10.5555/3001460.3001507

24. Rahman, N. H., Rainis, R., Noor, S. H., & Mohamad, S. M. S. (2016). *The buffering analysis to identify common geographical factors within the vicinity of severe injury related to motor vehicle crash in Malaysia*. *World Journal of Emergency Medicine, 7*(4), 278–284. https://doi.org/10.5847/wjem.j.1920-8642.2016.04.007

25. District Department of Transportation. (2024). *High Injury Network* [Data set]. Open Data DC. https://opendata.dc.gov/datasets/DCGIS::high-injury-network-1

26. Pasolini, G., Toppan, P., Toppan, A., Bandiera, R., Mirabella, M., Zabini, F., Bonata, D., & Andrisano, O. (2024). *Comprehensive assessment of context-adaptive street lighting: Technical aspects, economic insights, and measurements from large-scale, long-term implementations*. *Sensors, 24*(18), 5942. https://doi.org/10.3390/s24185942

27. Job, R. F. S. (2022). *Evaluations of speed camera interventions can deliver a wide range of outcomes: Causes and policy implications*. *Sustainability, 14*(3), 1765. https://doi.org/10.3390/su14031765

28. District Department of Transportation. (2022). *Vision Zero*. Government of the District of Columbia. https://visionzero.dc.gov/