TDM Validation Activity: Bicycle and Pedestrian Data

TETC Transportation Data Marketplace Data Validation

Prepared by: TETC Validation Team

9/4/2026

TDM Validation Activity: Bicycle and Pedestrian Data

TETC Transportation Data Marketplace Data Validation

Preliminary evaluation of bike/ped ancillary products sold through the Coalition's TDM.

TDM-VAL-10

Data Categories: Bike/Ped (Ancillary Product)

9/4/2026

The Eastern Transportation Coalition is a partnership of 20 states and the District of Columbia focused on connecting public agencies across modes of travel to increase safety and efficiency. Additional information on the Coalition, including other project reports, can be found on the Coalition's website: www.tetcoalition.org

Executive Summary

This report details the first bike and pedestrian ("bike/ped") data validation activity conducted by The Eastern Transportation Coalition (TETC) and focuses on evaluating commercial data products sold by five TDM vendors in the Transportation Data Marketplace (TDM): AirSage, INRIX, StreetLight, Replica, and LOCUS. To date, all prior validation activities have focused on Travel Time/Speed, Volume, or OD Datasets – all of which are well-defined core products in the TDM focusing on motorized travel. In contrast, this activity focuses specifically on active transportation travel, which includes vendor products that are less standardized, representing bike and pedestrian activity from diverse perspectives.

The goal of this validation study is to provide state DOTs with a detailed characterization of commercial bike/ped data products available in the TDM. Given the diverse characteristics of the datasets, the validation process is much less prescriptive than for prior studies – particularly those focused on Travel Time/Speed and Volume products. Rather than focusing on specific accuracy measures for well-defined products, the intent of this study is to provide agencies with sufficient context to understand the data landscape and make informed procurement decisions.

This validation activity is composed of two parts: (1) a review of each vendor's product offerings, and (2) a case study focusing on comparisons between vendor submissions and other available datasets. The analysis focuses on basic descriptive analysis of the datasets, sanity checks, and fundamental comparisons to available reference sources, including agency-managed bike/ped counters in Washington, DC and Arlington, VA, the American Community Survey, and sample data provided by Strava Metro. Several key takeaways emerged from the study:

- **Crowdsourced bike/ped products are diverse, making direct comparisons and evaluation challenging.** Although each of the five vendors offers data products that are categorized as "bike/ped", the datasets are distinct and quantify bike and pedestrian activity using different perspectives. For example, the bike/ped datasets vary in spatial resolution, with some vendors providing segment-level volumes (or normalized indices) and others providing zone-level trip estimates – either as trip volumes within an area or Origin-Destination flows or between geographies. Other differences include level of temporal aggregation, whether measures of activity represent samples of the population or population-level estimates. Additionally, particularly for zone-level results, reference/ground truth datasets are difficult – if not impossible – to acquire.

  Part I of this report provides a detailed characterization of all TDM bike/ped products to help provide clarity on each vendor's offerings.

- **Industry is in early product stage maturity but evolving quickly.** Five vendors in the TDM currently offer some type of scalable "bike/ped" product, and these offerings continue to evolve quickly. During the course of this validation activity, one vendor released an updated version of their model that substantially changed volume estimates relative to the initial version they had submitted. Another vendor announced the release of a new segment-level product that had not previously existed during the initial scan of TDM bike/ped products. This quickly changing landscape is encouraging, as it indicates that industry is working to bring better bike/ped data to market. However, when products are regularly updated, it can be difficult to use data for tracking trends over time, as changes in reported trips may be attributable to methodological differences rather than changes in user behavior.

- **Fundamental comparisons from the DC area case study yielded several insights at the zone and segment level:**

    - **Zone-Level (Bike):** Vendor datasets showed similar relative spatial patterns to one another and to a comparative dataset, although the magnitude of trip counts differed across vendors.
    - **Zone-Level (Pedestrian):** Vendor datasets generally identified central Washington, DC as a major concentration of pedestrian activity, but differed in total reported volume and in the spatial distribution of activity outside the urban core.
    - **Segment Level (Bike):** On a network scale, vendors showed varying spatial patterns, with only moderate agreement in which segments were associated with the highest bicycle activity. Follow-up analysis evaluated two vendors' bicycle volume estimates at agency count sites, providing an initial characterization of product accuracy on both roads and trails.

  Note that this case study focused specifically on the DC metro area, which has high levels of bike/ped activity. Further work is needed to understand whether the findings from DC area generalize to other locations with different land use, activity patterns, and data availability.

- **Benchmarks, metrics, and application-specific value need to evolve in parallel.** This study provides an initial characterization how well vendor data sources agree with one another and where possible, how closely they match reference datasets. These metrics and visuals serve as a starting point for quantifying vendor performance, but further work is needed to develop targets and benchmarks – a process that requires understanding agencies' intended use cases.

  For example, vendor products may be able to characterize spatial patterns – like identifying that Location A has twice the activity of Location B – well before they meet strict accuracy benchmarks for absolute volume. This ability to capture spatial patterns may be suitable for certain use cases (e.g., identifying high activity corridors, prioritizing projects, comparing exposure), but not others, reflecting the need to develop application-specific guidance.

Recommendations

In light of these findings, the validation team makes the following recommendations to agencies considering purchasing industry sourced bike/ped data products:

- **Make sure to understand product characteristics before purchasing.** Agencies are encouraged to review Part I of this report, which provides a summary and comparison of each vendor's product offerings. Because "bike/ped" data does not have a standard meaning, many diverse offers fall under the "bike/ped" label, even though some may not align with agency needs. For example, agencies should clarify whether they need zone or segment-level data and make sure that a candidate product supports the spatial and temporal granularity that is consistent with their intended applications.
- **Use industry sourced bike/ped data with caution.** Many of the industry-sourced bike/ped products included in this study appear promising and offer insights about bike or pedestrian activity in areas where no information currently exists.

  While this is exciting, *agencies should proceed with caution*, recognizing that many products are at an early product stage maturity and are evolving quickly. In particular, it is recommended that users:

- *Prioritize applications that require relative measures of bike/ped activity, rather than absolute trip volumes.* Initially, it is likely that these datasets will be more useful for applications such as infrastructure justification, project prioritization, safety/risk analysis, and other use cases where relative spatial patterns are sufficient to make decisions. While it is expected that absolute trip volume accuracy will only improve over time – and may be reasonable for filling gaps where no other option exists -- the accuracy of reported trip volumes cannot be verified at the zone level (no 'ground truth'), and segment-level accuracy is not yet fully understood, particularly for pedestrian datasets (not studied in this analysis) and for different geographies and land use contexts that are different than the metro DC area.
- *Avoid – or be careful – using industry-sourced bike/ped estimates for comparing bike/ped activity during different time periods*. It is difficult to know whether bike/ped metrics reported during two different time periods differ because of actual changes in bike/ped activity or other factors such as changes to a vendor's model/algorithm or differences in underlying data supply. Some vendors may have reasonable ways to control some of these factors (e.g., re-generate prior year estimates with an updated model), but even so, the uncertainty associated with bike/ped trip estimates may obscure changes and trends over time.

Additionally, the validation team recommends that the Coalition take several steps to help industry-sourced bike/ped products continue to mature:

- **Encourage standardization of bike/ped product offerings.** In order for the industry to focus and converge on usable and replicable data sets, it would be beneficial for the Coalition to define a set of standardized offerings which enable productive DOT functions for planning and operation, and can be procured from multiple vendors, compared, and validated. This should be considered a priority in the next procurement of the Transportation Data Marketplace (TDM V4).
- **Support agencies with traditional bike/ped count programs and encourage standardization of count data.** Traditional, infrastructure-based counts are critical sources of calibration for industry-sourced bike/ped products, so having more and better-formatted agency count data would help vendors improve their products and also aid with ongoing validation efforts. TETC recently developed the Active Transportation Count Specification (ATCS) to serve as a common data model and exchange format for collecting and sharing bike, pedestrian, and other non-motorized counts. It is recommended that the Coalition include this specification in the upcoming TDM V4 procurement and encourage its adoption by supporting agencies with its implementation.
- **Develop metrics and application-specific guidance**. In the same way that guidance has been developed for motorized traffic volumes, it would be useful for Coalition to provide a framework for interpreting the quality of bike/ped data. This would likely involve outreach across Coalition agencies to understand expectations for accuracy by use case and in different contexts and may also include characterizing the error associated with traditional methods for obtaining annual average bike/ped volumes. For example, following the approach taken by FHWA for motorized volumes, this could involve characterizing the errors associated with factoring a 48-hr or 1-week bike counts to an AADT-equivalent value, and setting targets for accuracy/precision that would be at least as good as the 'status quo' approach.

Table of Contents

Introduction

Transportation data sold through the Eastern Transportation Coalition (TETC) Transportation Data Marketplace (TDM) is procured from private industry based on contract specifications. The intent of the Coalition’s validation program has evolved from the original Vehicle Probe Project validation, which was primarily limited to ensuring that traffic data conformed to contractual standards for speed and travel time accuracy. Although the Coalition TDM still includes that essential function, it also has the flexibility to adjust to the needs of the Coalition members as the market evolves. The validation process is overseen by a technical advisory committee that sets general direction and reviews results.

The TDM validation includes both quantitative and qualitative analysis of datasets available through the marketplace as appropriate for each data type. The marketplace currently contains six core data items: Travel Time/Speed, Volume, Waypoint, Origin-Destination, Freight, and Conflation, with all but one (Travel Time/Speed) being sold through the marketplace for the first time. As such, the validation team, under the guidance of the TETC Validation Technical Advisory Committee (TAC), is establishing benchmarks and methods for effectively evaluating data quality and value applicable to the different data sets. Although validation activities are typically limited to core data areas, several Coalition agencies expressed interest in active transportation (“bike/ped”) datasets, so the TAC directed the validation team to evaluate the bike/ped data products available in the TDM.

Traditionally, measuring bike/ped activity has required deploying permanent or portable sensors to count pedestrians and cyclists at strategic locations. More recently, private industry has introduced novel data products that use Location Based Service (LBS) data from cell phones, and other sources to estimate bike and pedestrian activity at large spatial scales – analogous to similar products that have been developed for estimating motorized traffic volumes and origin-destination datasets. Several of these industry-sourced datasets are available as ancillary products in the TDM and are expected to become core data products in the next TDM procurement cycle.

This validation activity is composed of two parts: (1) a review of each vendor’s product offerings, and (2) a case study focusing on comparisons between vendor submissions and other available datasets, including agency-managed bike/ped counters in Washington, DC and Arlington, VA, and a survey dataset.

1. Data Vendors

Five vendors (either “primes” or subcontractors) who offer bike/ped data as an ancillary product in the TDM participated in the study: AirSage, INRIX, StreetLight, Replica, and LOCUS. Each vendor was highly engaged in the process, provided timely responses to questions and assisted with obtaining and interpreting the sample data for the analysis.

2. Product Characterization (Part I)

Although each of the five vendors offers scalable data products that are categorized as “bike/ped”, the datasets are distinct and quantify bike and pedestrian activity using different approaches and perspectives. For example, the bike/ped datasets vary in spatial resolution, with some vendors providing segment-level volumes (or normalized indices) and others providing zone-level trip estimates – either as trip volumes within an area or Origin-Destination flows or between geographies. Other differences include level of temporal aggregation, whether measures of activity represent samples of the population or population-level estimates, and the extent to which paths and trails that don’t follow the vehicle road network are considered.

These product differences present a similar problem to what was encountered for prior OD studies; unlike Travel Time/Speed and Volume products, for which all vendors deliver identical data with common formats and time aggregation, most of the bike/ped products are not directly comparable to one another. Furthermore, at least in the case of zone-level volumes and OD flows, reliable ‘ground truth’ data is difficult, if not impossible, to acquire. Thus, the main goal of this study is to characterize the available datasets.

As such, Part I of this report provides an overview of each of the vendor offerings, focusing on **product characteristics** that DOTs can use to compare data. Table I provides a concise summary of each product in a compact, comparable format, while the following subsections provide more verbose descriptions for each vendor. The following is a brief explanation of the information captured in Table 1:

- **Overview:** Includes product name and a brief description. Note that in some cases, vendors may have more than one distinct product that is considered ‘bike/ped’.
- **Characteristics:** Describes important characteristics of the dataset.
  - *Travel Modes*: Indicates the travel modes reported. Of particular interest is whether the vendors report activity for bikes, pedestrians, combined bike/ped users, or other micromobility modes.
  - *Reflects Population Estimates*: Indicates whether datasets are intended to reflect population-level behavior (as opposed to reporting sampled trips)
  - *Segment-level Reporting:* Indicates whether data is reported along road/path segments (as opposed to zone or point-level geographies).
  - *Zone-level Reporting:* Indicates whether data is reported at the zone level (e.g., census geographies, TAZs, or other polygons).
  - *Time Aggregation*: Focuses on the temporal granularity for which trips are summarized. Common examples include ‘Average Daily Trips’ for a study period, perhaps reported separately for weekdays and weekends.
  - *Delivery Frequency*: Described how often the data product is released (e.g., quarterly).
- **Data & Methods**: Describes the underlying input data sources and methodology.
  - *Data Sources*: Reports data sources that vendors use as inputs (e.g., Location-Based Service (LBS) data)
  - *Data Processing & Estimation*: Description of methodology used to process raw input data and, where relevant, estimate population-level results

Table 1 – Product Characteristics Matrix

Product Summary	StreetLight	Replica	INRIX		AirSage	LOCUS
Overview						
Product Name	Active Transportation Monitor (ATM) & Active Transportation Snapshot	Active Transportation Trips	VRU Products & Probe Counts	Ride Report	Pedestrian Activity Density & Ped Matrix	Person
Product description	ATM: Reports estimated annual bike and ped activity for census tracts. Snapshot (2026 release): Reports estimated annual bike and ped activity for segments	Reports estimated seasonal bike and pedestrian activity along segments and between OD zones.	Reports an index of bike/ped activity and 15-min probe counts at the segment level.	Reports observed trips taken using shared e-bikes & scooters along segments.	Reports estimated pedestrian activity at 'high density' locations and between OD zones.	Reports estimated bike/ped trips between OD zones and along segments (as of 2026 release).
Characteristics						
Travel modes	- Bike - Pedestrian	- Bike - Pedestrian	VRU Idx: Bike+Ped Counts: Bike, Ped	- E-Bikes - Scooters	Pedestrian	- Bike - Pedestrian
Reflects population estimates?	Yes	Yes	No	No	Yes	Yes
Segment-level reporting?	Yes (Snapshot product)	Yes	Yes	Yes	No	Yes
Includes paths/trails?	Yes	Yes	No	Yes	-	Yes
Segment type	OSM (simplified)	Custom links & OSM	OSM (granular)	Custom links	-	Overture Maps
Zone-level reporting?	Yes (ATM product)	Yes	Yes	No	Yes	Yes
Zone geographies	Census tract	- Census geographies - Custom definitions	Quad keys (VRU Idx)	-	- Census geographies - Custom definitions	- Census geographies - Custom definitions
Includes OD?	No	Yes	No	-	Yes (Ped Matrix)	Yes
Time aggregation for trip counts	Daily average	Daily average (weekday, weekend)	Custom (VRU Index) 15 min counts	Total counts (quarterly)	Daily average	Daily average (weekday, Sat, Sun)
Delivery frequency	Annually	Bi-annual (Fall & Spring products)	Custom	Quarterly	Custom	Quarterly
Methods						
Data sources	LBS, Next Gen NHTS, and other supplemental sources	LBS, consumer and resident data, land use, economic activity	LBS	GPS data from micromobility fleet operators	LBS	LBS and other supplemental sources
Data Processing & Estimation	See Section 3.1	See Section 3.2	See Section 3.3	See Section 3.3	See Section 3.4	See Section 3.5

3.1 StreetLight

StreetLight has two distinct bike/ped products: Active Transportation Monitor, which reports annual bicycle and pedestrian metrics at the census tract level, and Active Transportation Snapshot, which reports annual bike/ped volumes at the road segment level.

Active Transportation Monitor

StreetLight's Active Transportation Monitor product reports annual bicycle and pedestrian metrics at the census tract level, as well as historical bike/ped metrics at a segment level (2019 through April 2022). Active Transportation Monitor metrics can be obtained via interactive web platform and include volume estimates and mode share estimates for pedestrian, bicycle, and vehicle modes.

Characteristics

- **Travel Modes**: Bicycle and Pedestrian (reported separately).
- **Reflects population estimates?** Yes.
- **Includes recreational trips?** Yes.
- **Segment-level Reporting?** No. Segment-level data is not currently supported in the ATM, although historical data is available prior to April 2022. Current segment-level data is now reported in the separate Active Transportation Snapshot product, described below.
- **Zone-level Reporting?** Yes. Bike/ped activity is reported at the census tract level, reflecting trips that *start* within each zone.
- **Time Aggregation**: Annual Average Daily Traffic, reflecting a single average value across the entire year. This average value is reported separately for bike and ped modes.
- **Delivery Frequency**: Released annually.

Methodology

- **Data Sources**: LBS data, Next Gen NHTS, and other supplemental sources.
- **Data Processing & Estimation**: StreetLight's methodology for estimating these bike/ped metrics is documented in a publicly available white paper (link), and uses several data sources, including LBS trip data (representing a sample of trips across different modes), the Next Gen National Household Travel Survey (Next Gen NHTS) dataset, OpenStreetMap, and US Census geographies. Their estimation modeling process involves (1) determining modal split for metropolitan statistical areas (MSAs) using LBS and Next Gen NHTS data, (2) estimating total active transportation volumes at the MSA level using the modal splits and vehicle volume estimates, and (3) splitting the total active transportation volumes into bike and ped volumes at the tract level using LBS data.

Active Transportation Snapshot

StreetLight's Active Transportation Snapshot product contains segment-level estimates of both bike and pedestrian activity, which is a more spatially granular offering than their existing zone-level product (Active Transportation Monitor, described above).

Characteristics

- **Travel Modes**: Bicycle and Pedestrian (reported separately).
- **Includes recreational trips?** Yes.
- **Reflects population estimates?** Yes.
- **Segment-level Reporting?** Yes. StreetLight reports volume on a simplified version of OpenStreetMap segments. The OSM network is modified to combine parallel facilities, each of which may usually be represented as a separate feature. For example, a roadway with general purpose lanes, a cycle track, and two sidewalks may all be condensed into a single segment. Metrics include both directions of travel on a segment.
- **Zone-level Reporting?** No.
- **Time Aggregation**: Annual Average Daily Traffic, reflecting a single average value across the entire year. This average value is reported separately for bike and ped modes.
- **Delivery Frequency**: Released annually.

Methodology

- **Data Sources**: LBS data, 2000+ bike/ped counters, point of interest datasets, road/path network characteristics, and other contextual information.
- **Data Processing & Estimation**: StreetLight uses a common data processing method for ingesting input data, the details of which are outlined in a public whitepaper (link). At a high level, this approach includes ingesting and cleaning raw LBS data, defining trips, adding contextual information (e.g., road network, land use), and normalizing and expanding sample trips to estimated trip counts – typically leveraging machine learning techniques and observed counter data. This general data processing framework is applicable to all modes of travel, although the details of how sample trips are expanded to population-level estimates differ by mode.

  To develop bicycle and pedestrian volume estimates, they first use *vehicle* penetration rates in the target area to approximate bike penetration rates (taking advantage of the fact that both penetration rates follow similar distributions, but vehicle penetration rates can be more reliably measured). Afterwards, they scale the pedestrian and bicycle volume estimates to account for national differences in bicycle and vehicle penetration rate by training a regression model on observed counts.

Additional Notes

StreetLight provides guidance on suitable use cases for this segment-level product. They recommend that segment-level metrics be used for prioritizing active transportation focus areas, selecting corridors for facility/network improvements, and filling gaps in bike/ped metrics where there are no other data sources. However, at this point they caution against using it for before-and-after studies (external factors such as data supply and weather may impact volumes more than underlying changes in behavior) and other use cases where high precision of spatial granularity (e.g., sidewalk vs road) or temporal granularity is required.

3.2 Replica

Replica's Active Transportation Trips dataset estimates population-level bike and pedestrian activity (alongside other modes) using an activity-based model, resulting in highly granular measures of bike/ped activity that can be reported at the individual trip level (i.e., disaggregate trips with start/end points and timestamps), between Origin-Destination zones, or along specific road segments. Datasets reflect trips taken during a typical weekday and typical weekend in each season. Customers can interact with the data through a web interface, export results into common data and GIS formats, and query programmatically by accessing an online database.

Characteristics

- **Travel Modes**: Bicycle and Pedestrian (reported separately).
- **Includes recreational trips?** No. The methodology focuses on "purposeful non-motorized travel", and for example, would not include a walk around the block that starts and ends at the home location.
- **Reflects population estimates?** Yes.
- **Segment-level Reporting?** Yes. Replica reports volumes on its own custom, directional segments, and also associates each segment with an OpenStreetMap segment id.
- **Zone-level Reporting?** Yes. By default, Replica supports zone-level reporting for common census geographies (e.g., tract, block group) but can also support queries for user-defined geographies.
- **Time Aggregation**: Average Daily Traffic (seasonal), reported separately for weekdays vs weekends (also reported separately for bike/ped) over the seasonal period.
- **Delivery Frequency**: Released twice a year (Spring and Fall seasonal models).

Methodology

- **Data Sources**: Public Use Microdata Samples (PUMS), American Community Survey (ACS), built environment data, LBS data, point-of-interest datasets, and consumer transaction data.
- **Data Processing & Estimation**: Replica starts with a synthetic population ("replica") of everyone in the US census, models travel behavior in terms of activity sequences and locations and then determines mode choice for each trip by considering the route along different available modes. Large-scale simulations are used to apply these behavioral models to the synthetic population, ultimately resulting in mode-specific trips and associated volumes along each road segment. The outputs are then compared with local data – including 'reference' counts – to calibrate the model. More information about the methodology and data schema can be found here.

Additional Notes

Although Replica's estimation engine produces estimates for bike/ped volumes under existing conditions, its modeling approach is designed to also consider latent demand (i.e., travel demand that is not currently realized). Their modeling process assigns trips to specific travel modes (such as biking or walking) using a multi-modal route choice model that considers the available infrastructure. Using this framework, they can evaluate the impact of policy decisions that would change routing and mode choice outcomes (e.g., adding active transportation infrastructure that results in safe routing options for cyclists and pedestrians).

3.3 INRIX

INRIX offers two types of bike/ped products: a Vulnerable Road User (VRU) product, which focuses on probe-based bike/ped activity at the segment level, and Ride Report, which helps agencies monitor their shared micromobility (e-bike and scooter) fleet activity.

INRIX Vulnerable Road User (VRU) products & Probe Counts

INRIX's main active transportation product is called the VRU (Vulnerable Road User) Index, which is part of their Safety Analytics platform. The VRU Index, derived from probe data sources, describes the average proportion of non-vehicular road and street users (combined cyclists, pedestrians, and other users) on a road segment relative to other segments in a study area. It can either be presented as a heatmap (visually communicating the index for a region) or a value associated with specific road segments.

In addition to the VRU index, INRIX can provide raw probe counts for bike and pedestrian modes along each segment.

Characteristics

- **Travel Modes**:
    - VRU Index: Bike & Pedestrian modes are combined and referred to as "VRU".
    - Probe Counts: Bike & Pedestrian modes are reported separately.
- **Includes recreational trips?** Yes.
- **Reflects population estimates?** No.
    - VRU Index: Reflects *relative* volumes within a study area.
    - Probe Counts: Reflects observed samples, with no scaling to population level.
- **Segment-level Reporting?** Yes. INRIX uses a custom version of the OpenStreetMap network, breaking each OSM way (segment) into more granular subsegments. Both VRU indices and probe counts are reported on granular OSM segments, with values reported separately for each direction of travel.
    - Note: VRU metrics are only reported on segments that allow vehicle traffic and thus do not include trails/paths.
- **Zone-level Reporting?** INRIX uses color-coded quad-keys (a spatial indexing system) to visually communicate where there is more or less VRU activity. This is *not* a zone-level OD dataset.
    - Note: VRU metrics are not limited to roads with vehicle traffic, and thus can include trails, parks, campuses, and other facilities open to walking and biking.
- **Time Aggregation**:
    - VRU Index: Indices can be calculated for a custom-defined period.
    - Probe Counts: 15-minute aggregation for specific dates
- **Delivery Frequency**: Custom.

Methodology

- **Data Sources**: LBS data
- **Data Processing & Estimation**:
  - VRU Index: INRIX uses a machine learning model to classify each trip's mode (e.g., car, bike, ped, other) and then combines bike/ped trips into a combined 'VRU' category. These trips are then routed along the road network and associated with road segments, and the VRU index for each segment is calculated as the ratio between bike/ped trips on that segment and the average of all bike/ped trips for segments within a defined area of interest.
  - Probe Counts: The same initial process is used to classify each trip's mode and associate each observation with corresponding road segments. Afterwards, the probe counts are summarized for each segment and direction, without any extrapolation or inferences about population-level activity.

Ride Report

The Ride Report product (owned by INRIX), ingests and standardizes data from shared mobility operators (typically e-bikes and e-scooters), offering cities a unified and comprehensive view of micromobility usage in their jurisdictions. In contrast to many of the products mentioned here, this product does not estimate or model bike/ped activity. Instead, it ingests the data produced by e-bike/scooter operators, standardizes it, and presents segment-level trip counts and other statistics across areas of operation.

Characteristics

- **Travel Modes**: E-bike, scooter
- **Includes recreational trips?** Yes.
- **Reflects population estimates?** No.
- **Segment-level Reporting?** Yes. Ride Report uses a custom segment definition, with geometries based on OpenStreetMap. A single value is reported for each segment, reflecting both travel directions combined.
- **Zone-level Reporting?** No.
- **Time Aggregation**: No averaging – metrics reflect total observed trips recorded during selected period (usually 3 months)
- **Delivery Frequency**: Quarterly.

Methodology

- **Data Sources**: GPS data from shared mobility fleet operators.
- **Estimation Methodology**: The data processing steps are not provided in detail. However, at a high level, ingests the GPS data produced by e-bike/scooter operators, standardizes it, and presents segment-level trip counts and other statistics across areas of operation. No estimation/modeling is done to infer activity at the population level.

3.4 AirSage

AirSage offers two types of products, both focused on pedestrian activity: an OD matrix focusing on walking trips between zones and an Activity Density product that focuses on walking "hotspots".

Pedestrian OD Matrix

AirSage's Pedestrian OD matrix product estimates population-level pedestrian trips between census geographies (or other custom zone definitions). This data product is identical to their other Origin-Destination datasets but is filtered by mode to only reflect walking trips.

Characteristics

- **Travel Modes**: Pedestrian
- **Includes recreational trips?** Yes.
- **Reflects population estimates?** Yes.
- **Segment-level Reporting?** No.
- **Zone-level Reporting?** Yes. Supports census geographies as well as custom user-defined zones.
- **Time Aggregation**: Average Daily Trips, reported separately for weekdays and weekends during aggregation period (e.g., 1 month).
- **Delivery Frequency**: Custom.

Methods

- **Data Sources**: LBS data sourced from opted-in smartphone applications.
- **Data Processing & Estimation**: AirSage begins by preprocessing the raw GPS observations from its data panel – a process that involves removing duplicate sightings and applying other data cleansing techniques. Afterwards, it uses proprietary algorithms to analyze each device's movement patterns to infer home & work locations and assigns a "point type" to each GPS observation (e.g., transient point vs end point). The mode is inferred by a combination of movement patterns, with trip distance and travel speed used to identify pedestrian movements. Finally, these sample trips are extrapolated to population-level trip estimates. At a high level, they use devices' home locations to compute monthly weighted factors by census tract and then apply these weighted factors to each trip to obtain an estimate of all movements within the study area.

Pedestrian Activity Density

AirSage's Pedestrian Activity Density product identifies "hot spot" locations where many pedestrian movements are observed. It divides an area into a rectangular grid and reports the trip counts observed in each cell, identifying high-demand areas.

Characteristics

- **Travel Modes**: Pedestrian
- **Includes recreational trips?** Yes.
- **Reflects population estimates?** Yes.
- **Segment-level Reporting?** No.
- **Zone-level Reporting?** Yes. Trip counts are associated with rectangular grid cells. AirSage supports multiple grid resolutions, including 10-meter, 100-meter, and 1,000-meter cells, depending on the application and customer requirements.
- **Time Aggregation**: Average hourly trips, reported separately by day of week during aggregation period (e.g., 1 month).
- **Delivery Frequency**: Custom.

Methods

- **Data Sources**: LBS data sourced from opted-in smartphone applications.
- **Data Processing & Estimation**: AirSage uses the same data processing techniques as described above. However, for this application, they focus specifically on transient (moving) points that are considered walking trips, assign each point to the appropriate grid cell, and sum the estimated trips for each cell.

3.5 LOCUS

Overview

LOCUS's Person product estimates population-level bike and pedestrian activity between OD zones, and as of 2026, along road segments. Average daily trips are reported by mode for weekdays, Saturdays, and Sundays during each quarter. Additionally, Customers can interact with the data through a self-service web platform, or LOCUS can provide customized datasets for direct download.

Characteristics

- **Travel Modes**: Bicycle and Pedestrian (reported separately).
- **Includes recreational trips?** Yes.
- **Reflects population estimates?** Yes.
- **Segment-level Reporting?** Yes. LOCUS uses Overture Maps segments with OpenStreetMap tags (to facilitate crosslinking to OSM segments).
- **Zone-level Reporting?** Yes. By default, LOCUS supports zone-level reporting for common census geographies but can also support queries for user-defined geographies.
- **Time Aggregation**: Average Daily Traffic, reported separately by mode for weekdays, Saturdays, and Sundays for each quarterly period.
- **Delivery Frequency**: Released quarterly.

Methodology

- **Data Sources**: LBS data sourced from opted-in smartphone applications, as well as additional contextual sources (population, employment, and school enrollment; built environment and land use; road/path network characteristics, environmental conditions). Travel surveys are also used for validation.
- **Data Processing & Estimation**: At a high level, LOCUS's process begins by identifying activity locations and trips. Raw LBS traces are processed to identify major activity locations (e.g., home, regular locations) and to determine movements between them. Afterwards, each trip is classified by purpose (home–regular, home–other, etc.) and assigned a mode (walk, bike, motorized). Mode labeling uses a behavioral approach that incorporates trip characteristics, network context, land-use patterns, and environmental factors such as weather and terrain. Finally, device-level data are expanded using demographic and activity-based controls so that the resulting outputs represent average daily person travel for the full population. Additional adjustments normalize trip-making across devices and correct for under representation of short trips. Final estimates are evaluated against national and regional travel survey benchmarks and a set of behavioral and geographic reasonableness checks. Segment level volumes are also validated against available active transportation counts, where available.

3. Case Study (Part II)

Given that "bike/ped data" encompasses many different types of products with varying spatial/temporal resolutions and assumptions (see Part I), it is difficult to statistically compare such datasets to independent sources of "ground truth" to compute statistical accuracy measures. In some cases (e.g., zone-level census tracts or block group data) there is no independent "ground truth" dataset against which vendor data can be compared. At the segment level, reference counts – usually collected by agencies – are quite sparse and often used by the vendors in calibration.

As such, this case study focuses primarily on gathering sample data from each vendor for a common geography (and where possible, the same time period) and performing several fundamental comparisons between datasets. These fundamental comparisons were selected by the TDM Technical Advisory Committee and intended to capture the most relevant, straightforward analyses at both the Segment and Zone levels.

4.1 Study Parameters

Study Area

The DC metro area – defined as the Washington-Arlington-Alexandria Metropolitan Statistical Area (MSA) – was selected as the study area for this activity, based in large part on District DOT's interest in bike/ped data and the availability of agency count data from DDOT (DC) and the City of Arlington (VA). Vendors were given a GIS file with the CBSA boundary for reference but told that the analysis would focus on three counties/areas in particular: DC, Arlington (VA), and Prince George's County (MD). In the end, the focus area ended up being limited to DC and Arlington, as shown in Figure 1.

<u>Note</u>: Vendor data quality in the DC metro area may not necessarily be representative of all other cities, geographies, and land use contexts. However, the prevalence of walking and biking in the DC area – plus the availability of various comparative datasets – makes this a useful first location to characterize vendor performance.

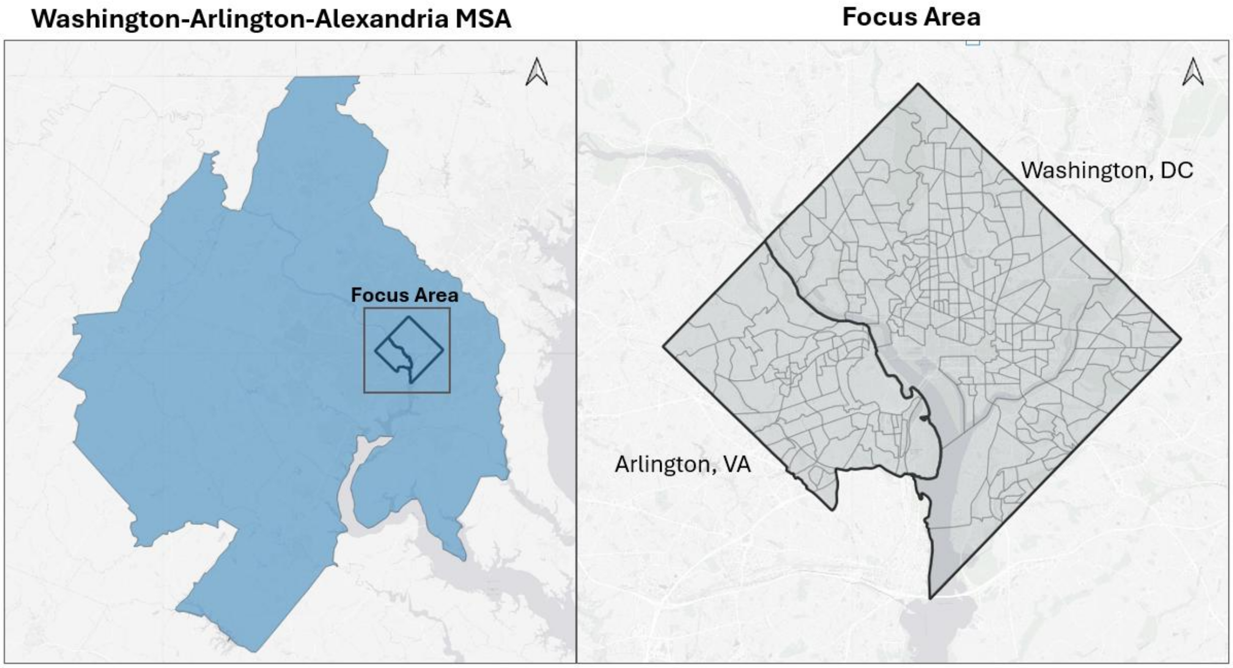

Figure 1 – Study area focusing on subset of Washington-Arlington-Alexandria MSA

Time Period

Vendors were asked to provide the most recent data available, targeting calendar year 2024 for annual products, and Fall 2024 for seasonal products, where possible.

Other Parameters

The data request to vendors was open-ended and much less prescriptive than typical TDM validation activities. Recognizing that vendors produce data with different temporal aggregations (e.g., annual products vs quarterly) and for different spatial definitions (e.g., segments vs zones), the validation team accepted each vendor's submissions without many restrictions. The key parameters were:

- **Segment level** datasets must include geometry definitions for each segment.
- **Zone level** datasets must be reported at the Census Tract level (or a more granular level, such as Census Block Groups, which can be converted to tracts).

Datasets

Zone-Level

Table 2 summarizes the vendor products submitted at the Zone level for this case study. Three of the vendors submitted OD flows between Census Block Groups (BG), while one submitted trips originating from Census Tracts. Because BGs can be aggregated to Tracts and OD flows can be aggregated to total trip origins, the lowest common denominator for all vendors is trip counts originating from Tracts. There are three vendors who submitted metrics for bike trips (StreetLight, Replica, LOCUS) and a slightly different group of three who submitted metrics for pedestrian trips (StreetLight, AirSage, LOCUS). All submissions focus on population-level estimates (rather than sample counts) and report the average daily trips, although there are some differences in the time period used for aggregation, and whether results are aggregated across all days of the week or sub-reported for different day types.

Table 2 – Zone-level vendor submissions

Vendor	Streetlight	Replica	AirSage	LOCUS
Zone Type	Census Tract (Origins)	Census Block Group (O-D)	Census Block Group (O-D)	Census Block Group (O-D)
Time Period	2023 (annual)	Fall 2024 (3 months)	Oct. 2024 (1 month)	Q3 2024 (3 months)
Day Type	All	Weekday	Weekday, Weekend	Weekday, Saturday, Sunday
Mode	Bike, Pedestrian (separate)	Bike	Pedestrian	Bike, Pedestrian (separate)
Metric Type	Population Est.	Population Est.	Population Est.	Population Est.
Metric Aggregation	Daily avg.	Daily avg.	Daily avg.	Daily Avg.

Comparative Datasets

At the zone level (particularly for Census Tracts), it is very difficult to establish "ground truth" measurements for bike or pedestrian activity. However, two datasets – referred to as "comparative datasets" rather than "ground truth" – were identified as opportunities to sanity check patterns reported by the vendors.

- **5Y American Community Survey (2019-2023)**

  The 2023 5-year American Community Survey (ACS) was used as a comparative dataset for tract-level bicycle and pedestrian activity. Specifically, ACS journey-to-work data was used to identify resident workers whose primary commute mode was bicycle or walking. These values were summarized by Census Tract and used as a residence-based indicator of bicycle and walking commute origins. These ACS values are population estimate – similar in that respect to the vendor datasets, rather than raw sample counts.

  ACS provides useful information because it is publicly available, consistently reported across Census Tracts, and includes both bicycle and walking commute modes. However, *it has important limitations for this validation activity*. ACS captures only commute travel and only the primary mode used for the journey to work. It does not include non-work trips, recreational trips, school trips, access trips to transit, visitor activity, or other forms of bicycle and pedestrian travel. In addition, ACS is residence-based, meaning that trips are associated with where workers live rather than where total bicycle or pedestrian activity occurs. As a result, ACS may underrepresent activity in major employment, commercial, recreational, or institutional areas with limited residential population. For these reasons, *ACS is used only as a broad comparative reference and should not be interpreted as a ground truth measure of total bicycle or pedestrian trips*.

- **Capital Bikeshare Trips Dataset**

  Capital Bikeshare (CaBi) is a publicly owned, docked (point-to-point) bike sharing system that serves the DC metropolitan area. It has about 8,000 bikes and over 700 stations across the DC metro area, with coverage in DC and Arlington, VA.

  CaBi makes monthly trips datasets available to the public, which includes information about start/end station, timestamp, and other additional trip statistics. Although this is not a zone-level dataset, the station-to-station trips can be aggregated to the zone level (by assigning each station to the appropriate census geography) and summarized at the same spatial and temporal scale as the vendor datasets. Note that this approach is only suitable in areas with a high density of stations, such as DC and Arlington.

  For this analysis, CaBi trip origins were aggregated to the Census Tract level and converted to average daily bicycle trip origins for comparison with the vendor bicycle datasets. Average daily trip origins were calculated using weekday trips during fall 2024, because this time period and day type most closely align with the three vendor datasets providing bicycle trip estimates. This provides a useful independent point of comparison in areas with dense bikeshare station coverage, particularly central Washington, DC and parts of Arlington. However, CaBi represents only bikeshare activity and does not capture private bicycle trips or bicycle trips occurring outside the bikeshare service area. Its usefulness as a comparative dataset is therefore strongest in areas with high station density and weaker in areas where bikeshare access is limited or unavailable. As with ACS, CaBi is used to help interpret spatial patterns in bicycle activity, not as a comprehensive measure of total bicycle trips.

  Figure 2 shows the 508 Capital Bikeshare stations used in this analysis within Washington, DC and Arlington, Virginia, overlaid on Census Tract boundaries. The map illustrates the spatial coverage of the CaBi stations that were used to assign trip origins to tracts. Station coverage is densest in central Washington, DC and extends into several parts of Arlington, which supports the use of CaBi as a comparative dataset in these areas while also highlighting that the comparison is most meaningful where station availability is relatively dense.

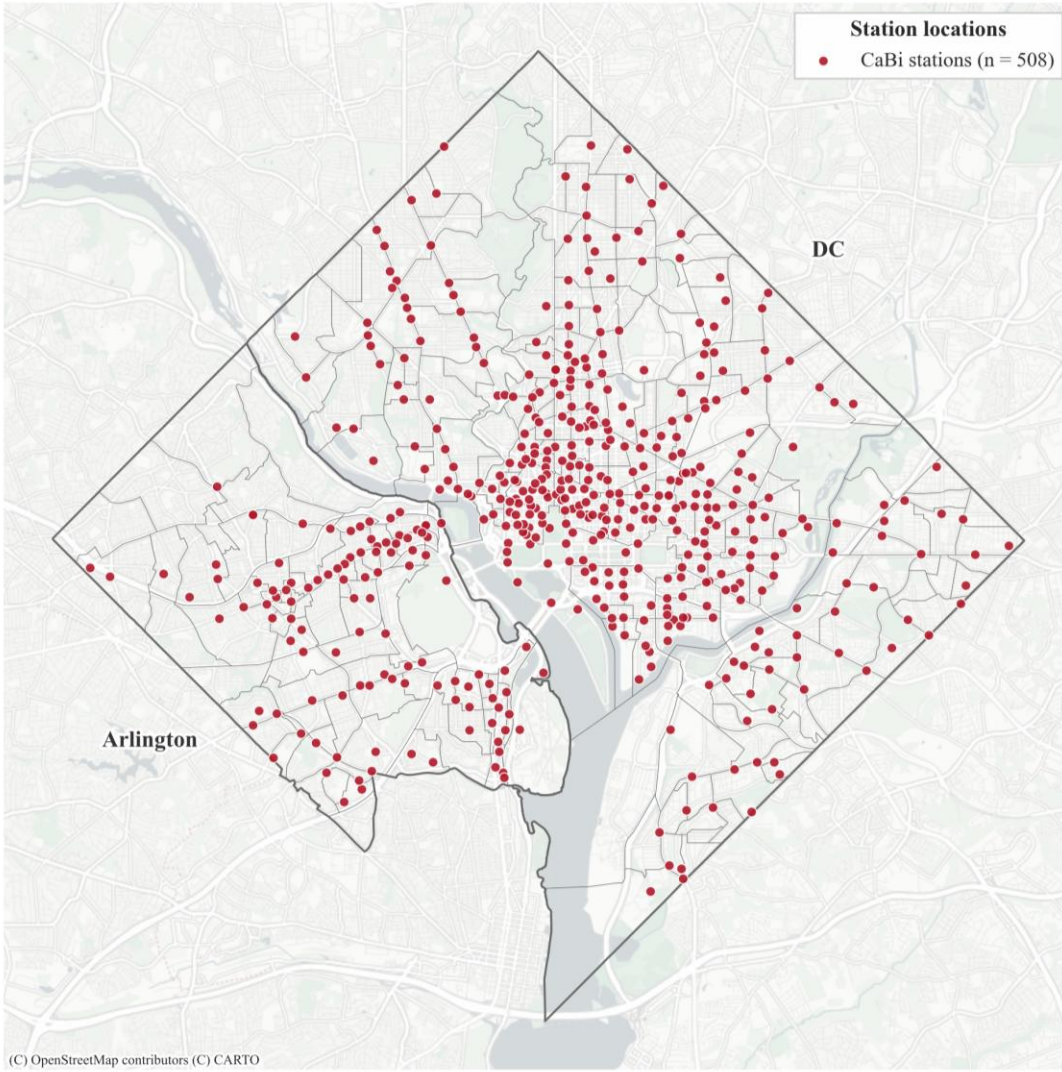

Figure 2 – Capital Bikeshare station coverage in Washington, DC and Arlington, Virginia

Segment-Level

Table 3 summarizes the vendor products submitted at the Segment level. StreetLight and Replica are most similar, with both reporting segment-level estimates that reflect population-level bike or pedestrian activity. INRIX's VRU product is a bit different, reporting segment-level indices that capture relative volume of bike/ped activity as compared to other roads in the study area. INRIX's Ride Report product reports e-bike trips that are part of the city's shared mobility fleet. These products use different time periods and differ in whether results are aggregated across all days of the week or sub-reported for different day types.

Note that INRIX also submitted bike and pedestrian probe counts (reflecting observed samples, not population estimates) for much more granular 15-minute periods during a single day in July 2025. Given that this data set is unlike the others and for a very limited period, it is not included in the analysis or Table 3.

Table 3 – Segment-level vendor submissions

Vendor	Streetlight	Replica	INRIX (VRU)	INRIX (Ride Report)
Segment Type	**Custom** Simplified version of OSM	**Custom** Linked to OSM	**OSM** OSM segments divided for increased granularity	**Custom** Custom link definitions
Travel Direction	Both directions combined	Directional	Directional	Both directions combined
Mode	Bike, Pedestrian (separate)	Bike	Bike & Pedestrian (combined)	E-bike
Time Period	2024 (annual)	Fall 2024 (3 months)	May 2025 (1 month)	Q3 2024 (3 months)
Day Type	All	Weekday, Weekend	All	All
Metric Type	Population Est.	Population Est.	Index (relative volume)	Observed samples
Metric Aggregation	Daily avg.	Daily avg.	Direct value	Total counts

Comparative Datasets

At the segment level, the following data sources are used for comparative purposes.

- **Agency counts**

  Agency bicycle and pedestrian count data collected by the District Department of Transportation (DDOT) and Arlington County, Virginia were made available through the BikePed Portal. These data were used specifically for the segment-level evaluation at agency count sites discussed later in this report, where vendor-reported segment volumes are compared with observed counts at matched locations.

  For the bicycle count-site evaluation, the final set included 68 counters, consisting of 49 trail counters and 19 road counters. Five additional count sites in Montgomery County, Maryland, corresponding to 10 directional counters, were included outside the core DC-Arlington study area to increase the number and diversity of comparison locations. The utilized bicycle count sites are shown in Figure 3 by facility type.

The agency counts provide observed volumes at specific monitored facilities, such as a trail, bicycle lane, or roadway location. They do not represent a complete network-wide ground truth dataset. In addition, count locations do not always align directly with vendor segment definitions, particularly because vendors differ in whether segments are directional or bidirectional and how trails, roadways, and bicycle facilities are represented. Therefore, a manual crosswalk was created to match each usable count location to the most appropriate vendor segment or segments. The count data were then processed separately for each vendor comparison to account for differences in time period, directionality, and reporting format.

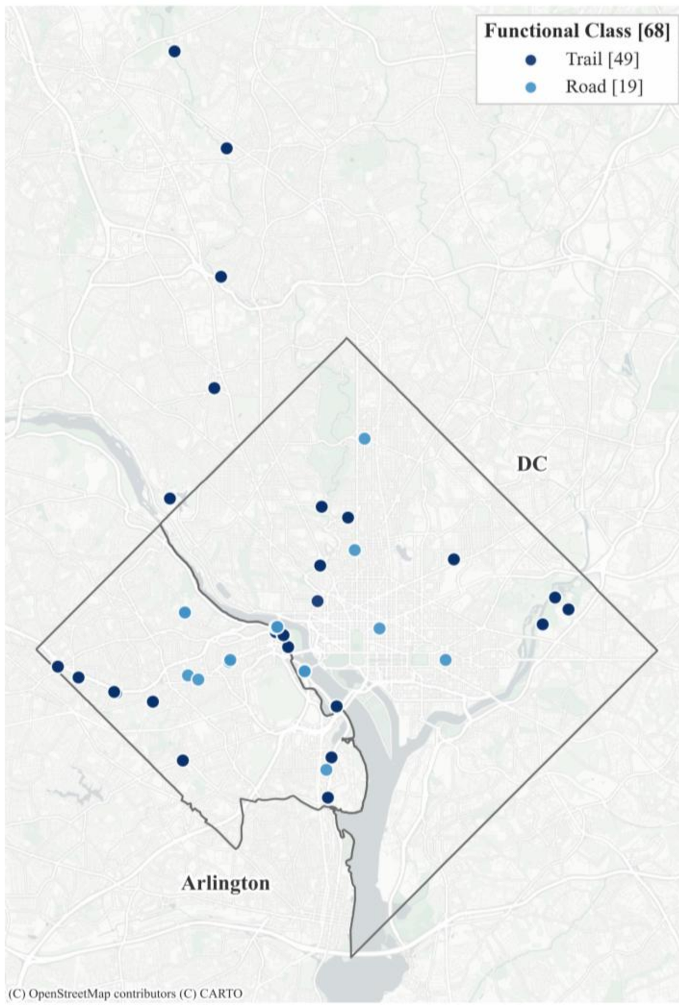

Figure 3 – Agency bicycle count sites used in the segment-level evaluation

<u>Note</u>*:* DDOT and Arlington both have web interfaces where data from agency counters can be viewed publicly, meaning that the reference data for this validation cannot be considered truly 'blind'. However, these agency dashboards do not provide the exact volumes used for validation; they do not always report counts at the same level of detail (e.g., travel mode and direction), and focus on raw volumes reported by counters before any quality control routines or correction factors are applied. The final count data – used to evaluate vendors – was acquired through BikePed Portal, which implements its own

data quality processes and calculates official volume aggregations. *BikePed Portal and its underlying database was not made available to vendors.*

- **Strava Metro**: Strava Metro is an aggregated dataset that summarizes biking and walking activity recorded by users of Strava (a fitness and activity tracking app). When Strava members record an activity (e.g., bike ride, walk, hike), the details of their trip are incorporated into an aggregated dataset, with personally identifying information removed to protect privacy. This aggregated dataset can be used to provide insights into where people walk and ride.

  Strava Metro makes their data available through an online platform, where users can visually explore bicycle and pedestrian activity from different perspectives (e.g., origin-destination flows, hexagon counts, and edge counts) and export raw data for further analysis. For our purposes, we downloaded monthly edge counts (i.e., segment volumes representing observed sample trips from Strava users) in 2024 for bicycle activity, which encompasses both traditional and e-bike trips.

4.2 Data Processing and Standardization

A core challenge was wrangling the various vendor data and comparative datasets – each with different spatial scales, temporal resolutions, and reporting assumptions – into common data structures to facilitate meaningful comparisons. Rather than adapting our analysis to each vendor's custom format, we developed a schema and organized all datasets – as well as underlying spatial layers – into a PostGIS spatial database. This relational database includes not just the volume metrics, but also the underlying spatial information associated with zone, segment, and point locations (e.g., zone geographies, road segments, and count sites), and maintains data integrity through foreign keys and other constraints.

Specialized scripts were used to transform datasets – a mixture of CSV, Shapefile, GeoJSON, GeoPackage and other formats – into data structures that could be loaded into the database. This Extract, Transform, Load (ETL) process handled unit conversions, spatial projections, and explicitly encoded assumptions about how the metrics should be interpreted (e.g., probe samples or population estimates). Once the 'heavy lifting' is done up front to standardize the datasets, the database can be queried transparently (in scripts, GIS applications, etc.) to support subsequent analysis.

The specifics of the database implementation, and ETL processing are outside the scope of this report, but the validation team can share additional details with the interested reader, if desired.

4.3 Analysis Methods

The analysis methods were developed in coordination with the TDM Technical Advisory Committee (TAC) and were designed to reflect the exploratory nature of this bicycle and pedestrian data validation activity. Unlike prior TDM validation studies focused on standardized travel time, speed, or volume products, the active transportation datasets evaluated in this study differ substantially in spatial resolution, temporal aggregation, network representation, and whether the reported metrics represent population estimates, sample counts, or indices. As a result, the analysis focuses on descriptive comparisons, reasonableness checks, and limited comparisons to available reference or comparative datasets, rather than applying a single standardized accuracy framework across all products.

The analysis is organized into three parts, shown below in Figure 4: (1) Zone-level comparison of bicycle and pedestrian trips by Census Tract, (2) network-wide segment-level comparison of bicycle activity, and (3) segment-level evaluation at agency count sites.

These components provide complementary perspectives on the datasets by comparing activity patterns across common zones, examining high-activity corridors across segment-based products, and evaluating vendor estimates at locations where observed count data are available.

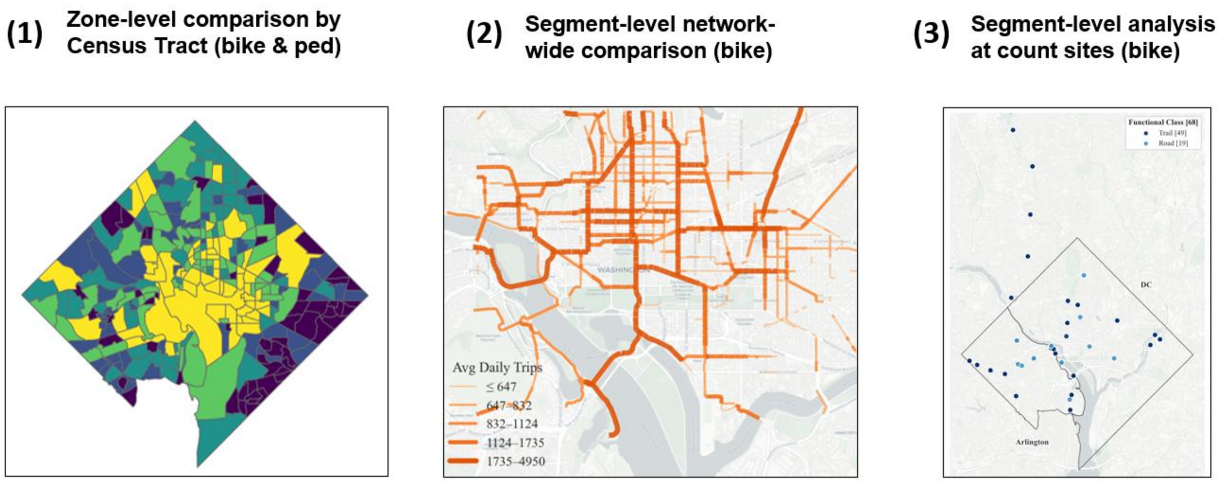

Figure 4 – Overview of analysis methods

Analysis #1: Zone-Level Comparison of Trips by Census Tract

The zone-level analysis used Census Tracts as the common geography for comparing vendor-reported bike and pedestrian activity (separately), with vendor datasets standardized to average daily trips originating from each Census Tract. Bicycle activity was compared across StreetLight, Replica, and LOCUS, with ACS commuting data and Capital Bikeshare trips included as comparative datasets. Pedestrian activity was compared across StreetLight, AirSage, and LOCUS, with ACS walking commute data included as a comparative dataset.

Because no comprehensive ground truth dataset exists for total bike or pedestrian trips at the Census Tract level, this analysis focused on comparisons across datasets. The primary emphasis was on comparing relative spatial patterns, overall magnitude, and the extent to which vendors identify similar concentrations of activity across the study area.

Analysis #2: Segment-Level Comparison (Network-Wide)

The network-wide segment-level analysis focused on bicycle activity reported along transportation network segments. This analysis included StreetLight and Replica as the two vendor datasets with segment-level bicycle volume estimates, with Strava Metro and Ride Report included as comparative datasets. The purpose of this analysis was to examine how each dataset represents bicycle activity across the broader network and whether the datasets identify similar high-activity corridors and locations.

In addition, INRIX's Vulnerable Road User (VRU) Index was evaluated separately as a segment-level measure of relative active transportation activity. Unlike the other datasets, the VRU Index does not estimate bicycle volumes and instead represents the relative concentration of combined bicycle and pedestrian activity on a roadway segment compared with other segments in the study area. Because the VRU Index combines multiple active transportation modes, excludes off-road trails and paths, and reports relative rather than absolute activity levels, it is not directly comparable to the bicycle volume datasets. Therefore, the VRU analysis focuses on identifying the roadway corridors that INRIX highlights as having elevated active transportation activity and comparing those patterns qualitatively with the other segment-level datasets.

A key methodological challenge is that each dataset reports activity on its own segment network. These networks are not always directly compatible for alignment because vendors may represent the same corridor differently. StreetLight generally reports activity on simplified bidirectional segments that may collapse multiple nearby facilities into a single roadway-aligned segment, while Replica reports activity on a more detailed directional network, with separate segments for individual facilities where applicable. Figure 5 illustrates this issue conceptually.

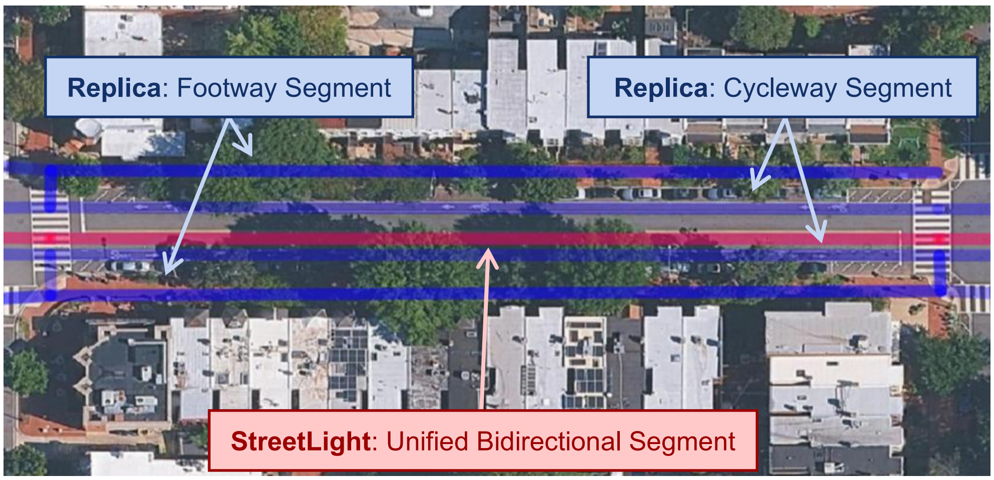

Figure 5 – Example of differences in vendor segment network definitions

Because of these differences, direct segment-to-segment conflation was not used as the primary network-wide comparison approach. Instead, the analysis relied on distributional summaries and visual comparisons. Segment-level bicycle activity was standardized to the extent possible, including converting directional estimates to bidirectional values where needed. High-activity segments were then identified within each dataset using relative thresholds, allowing the analysis to compare broader spatial patterns and corridor-level concentrations of bicycle activity across datasets without assuming that all vendor segments represent identical physical facilities.

Analysis #3: Segment-Level Evaluation at Agency Count Sites

The segment-level count-site evaluation compared vendor-reported bicycle estimates with agency-managed bicycle counts at selected locations. This analysis was limited to StreetLight and Replica, which both submitted segment-level bicycle volume estimates. Agency bicycle counts from DDOT, Arlington County, and other available count locations were used as the reference dataset.

Vendor segments were manually associated with relevant count locations where feasible, recognizing that vendor network definitions do not always align directly with the specific facilities measured by agency counters. Comparisons were conducted using correlation and error-based measures to evaluate agreement between vendor estimates and observed counts at these locations. This analysis provides a targeted count-site comparison, while the broader zone-level and network-wide analyses provide complementary perspectives on the vendor datasets.

4.4 Results and Discussion

Analysis #1: Zone-Level Comparison of Trips by Census Tract

The first set of analyses compares vendor-reported bicycle and pedestrian activity at the Census Tract level. As described in Section 4.3, this analysis uses trip origins as the common basis for comparison, with all datasets aggregated or summarized to average daily trips originating from each tract where possible. Because there is no comprehensive ground truth dataset for total bicycle or pedestrian trips at the tract level, the results are interpreted as cross-dataset comparisons rather than formal accuracy assessments. The analysis focuses on differences in overall trip magnitude, county-level totals, spatial distribution, and relative agreement across datasets.

Bicycle Trips

Among the vendor datasets, Replica reports the highest total bicycle activity, with approximately 113,500 average daily trips, followed by StreetLight (59,500) and LOCUS (47,500). In contrast, ACS and Capital Bikeshare report much lower totals, approximately 14,100 and 13,400 trips, respectively. These differences are expected because ACS reflects bicycle commuting only, while Capital Bikeshare captures only trips made through the regional bikeshare system.

Figure 6 shows the county-level distribution of the same totals for the District of Columbia and Arlington County. Across all datasets, most bicycle activity is concentrated in the District of Columbia, with substantially lower volumes in Arlington. The relative ordering of the vendor datasets is consistent in both counties: Replica reports the highest volumes, followed by StreetLight and LOCUS. These results indicate substantial differences in estimated trip magnitude across vendors, while also showing broadly consistent county-level patterns.

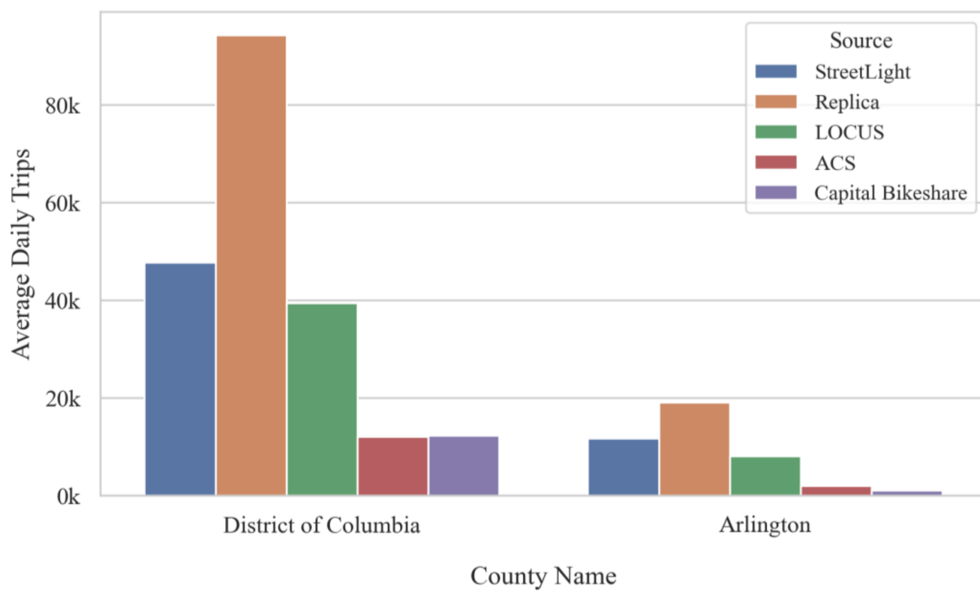

Figure 6 – Comparison of average daily bicycle trips by county

Figure 7 compares the spatial distribution of average daily bicycle trips originating from Census Tracts using dataset-specific percentile classes. For each dataset, Census Tracts were organized into five successive groups, starting with lowest reported bicycle activity to the highest reported bicycle activity. This was accomplished by dividing census tracts into five percentile categories based on the volume of trip origins within that dataset: 0-20th percentile, 20th-40th percentile, 40th-60th percentile, 60th-80th percentile, and 80th-100th percentile. The 0-20th percentile range, referred to as the 1st quintile, represents the census tracts with the lowest reported bicycle activity, and the 80th-100th percentile range, referred as the 5th quintile, represents the census tracts with the highest bicycle activity. Using this method, the maps in Figure 7 (color coded by quintile) emphasize relative spatial patterns rather than differences in absolute trip magnitude. This approach allows areas with comparatively high bicycle activity within each dataset to be identified and compared across data sources.

The color-coded quintile maps reflecting relative trip making behavior provide a more comparable view of spatial patterns across datasets. When each dataset is color-coded by quintile range distribution, StreetLight, Replica, and LOCUS show broadly similar concentrations of bicycle activity, with the highest-activity tracts clustered in the core of Washington, DC and adjacent high-activity areas. Capital Bikeshare also shows strong agreement with the vendor datasets in the core areas where station coverage is dense, suggesting that it is useful as a comparative dataset for identifying major bicycle activity concentrations. ACS shows weaker spatial agreement with the vendor datasets and Capital Bikeshare. This is expected because ACS captures only bicycle commuting and is based on the residential location of survey respondents, rather than the location where all bicycle trips originate. As a result, ACS does not fully capture bicycle activity in areas such as the National Mall and other core DC activity centers where residential population is limited and reported commute origins are sparse or unavailable.

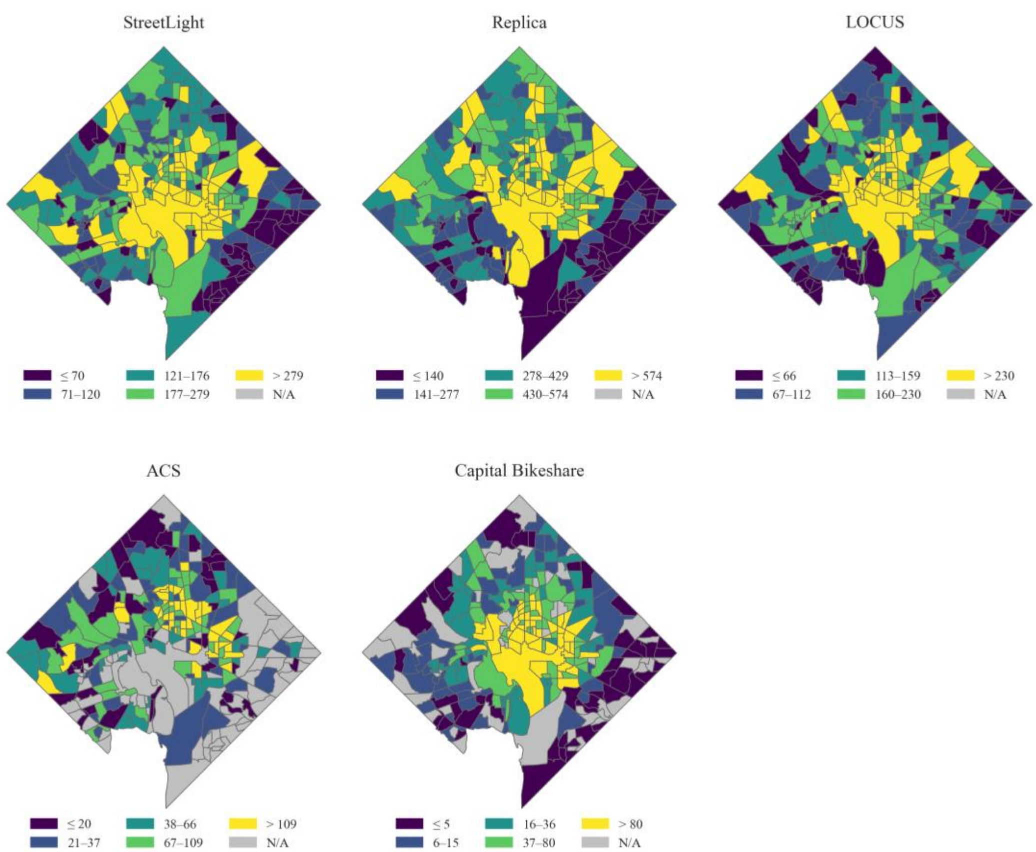

Figure 7 – Relative bicycle activity by census tract using dataset-specific percentile classes

As an additional sensitivity check, average daily bicycle trip origins were normalized by Census tract population and compared with the un-normalized results. The per-capita maps showed generally similar spatial patterns across the vendor and comparative datasets, with the highest relative activity remaining concentrated in central Washington, DC and nearby high-activity tracts. Normalization changed the classification of some individual tracts but did not materially alter the overall geographic patterns or the broader conclusions of the comparison. Census tracts with no reported residential population could not be assigned a per-capita value and were therefore excluded from the normalized analysis. Because the normalized results provided limited additional insight beyond the trip-volume maps, they are not presented separately in this report.

To further evaluate whether datasets identify similar high-activity areas, the top 10 percent of Census Tracts by reported bicycle activity were selected within each dataset and compared across sources. This analysis complements the choropleth maps, as shown in Figure 7, by quantifying overlap in the highest-activity tracts, while the subsequent correlation analysis evaluates agreement across the full set of tracts. Figure 8 highlights the top 10 percent of Census Tracts based on average daily bicycle trips reported by each data source. The maps show that the three vendor datasets generally identify similar high-activity origins, particularly in the core urban areas of Washington, DC. Capital Bikeshare shows a similar but more geographically concentrated pattern, reflecting the areas where bikeshare station coverage is strongest. In contrast, ACS highlights a smaller and more dispersed set of high-activity tracts, reflecting its narrower focus on bicycle commuting and the limited number of tracts with available bicycle commute estimates.

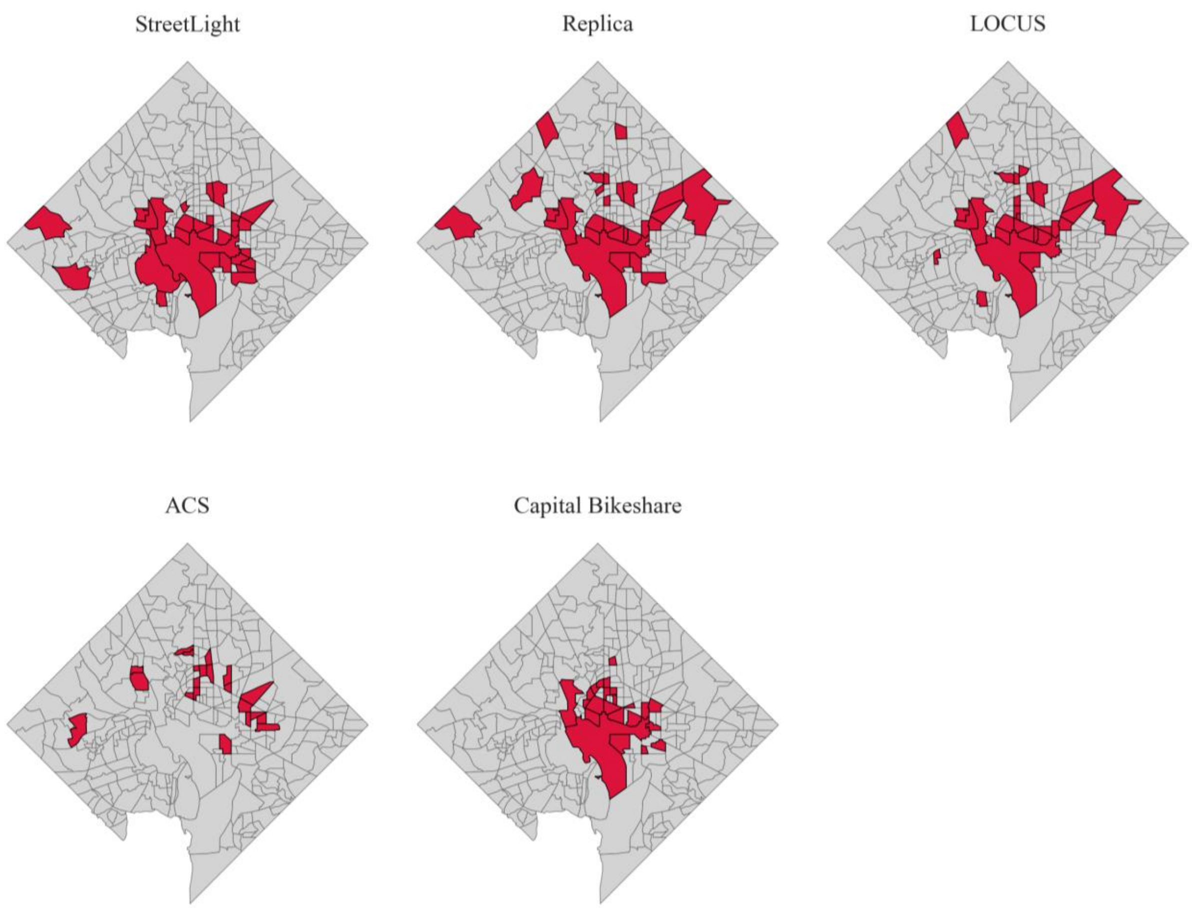

Figure 8 – Top 10% census tracts by average daily bicycle trips

The high-activity tract maps provide a qualitative visual comparison of where each dataset identifies the strongest concentrations of bicycle activity. To assess agreement more systematically, Spearman correlation coefficients were calculated across all Census Tracts. This correlation analysis evaluates whether datasets rank tracts similarly across the full range of reported bicycle activity, rather than focusing only on the highest-activity locations shown in the maps.

Figure 9 shows generally strong rank-order agreement among the three vendor datasets. StreetLight and Replica have the highest correlation (0.78), followed by Replica and LOCUS (0.75) and StreetLight and LOCUS (0.72). This indicates that, **although the vendors report different total bicycle trip volumes, they generally assign higher and lower activity levels to similar tracts**. Capital Bikeshare is also strongly correlated with the vendor datasets, with coefficients ranging from 0.67 with LOCUS to 0.73 with StreetLight. This further supports its use as a comparative dataset for identifying relative bicycle activity patterns in areas with sufficient bikeshare coverage. ACS shows weaker correlations with all other datasets, ranging from 0.23 with LOCUS to 0.39 with Capital Bikeshare. This is consistent with the previous maps, and reflects the fact that ACS captures only bicycle commuting rather than total bicycle activity.

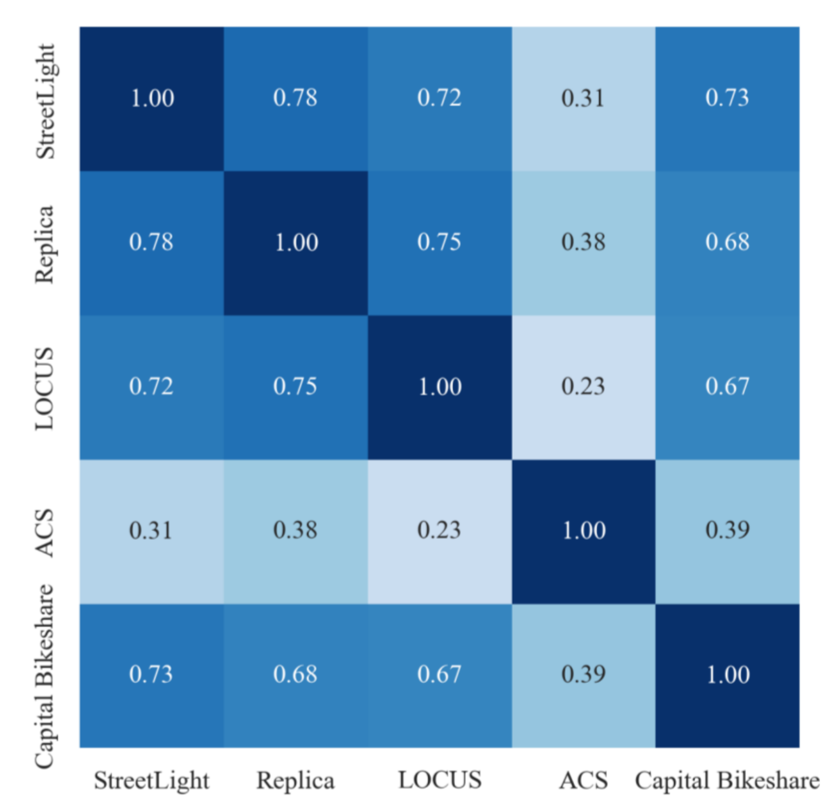

Figure 9 – Spearman correlation of tract-level average daily bicycle trips across data sources

As an additional sensitivity check, Spearman correlations were also calculated using average daily bicycle trip origins normalized by Census tract population. The population-normalized results were generally consistent with correlations based on total trip volumes. Rank-order agreement among the three vendor datasets remained strong, with coefficients ranging from 0.71 to 0.77, while ACS continued to show substantially weaker agreement with the other datasets. Correlations between Capital Bikeshare and the vendor datasets increased moderately after normalization, ranging from 0.76 to 0.77, but the overall interpretation remained unchanged: the vendor datasets and Capital Bikeshare identify broadly similar spatial patterns of bicycle activity, whereas ACS reflects a more distinct pattern. Because normalization did not materially change the conclusions, the additional correlation matrix is not presented separately.

Overall, the bicycle trip-origin analysis shows that the vendor datasets differ substantially in total estimated volume but are more consistent in their relative spatial patterns. StreetLight, Replica, and LOCUS generally identify similar high-activity tracts and rank tract-level bicycle activity in comparable ways, while ACS provides a more limited comparison point because of its narrower commute-based definition.

Pedestrian Trips

Among the vendor datasets, LOCUS reports the highest total pedestrian activity, with approximately 436,000 average daily trips, followed by StreetLight (360,400) and AirSage (221,000). ACS reports substantially fewer trips, approximately 86,700, which is expected because ACS captures walking-to-work trips only rather than all pedestrian travel. Although AirSage reports the lowest total among the vendors, its county-level distribution differs from the other two vendor datasets.

Figure 10 shows that StreetLight and LOCUS both report substantially higher pedestrian activity in the District of Columbia than in Arlington, while AirSage reports a more balanced distribution between the two jurisdictions. In Arlington, AirSage reports slightly higher pedestrian activity than both StreetLight and LOCUS, despite reporting a lower total volume overall. This suggests that the difference among vendors is not limited to overall scaling; the datasets also differ in how pedestrian activity is distributed spatially across jurisdictions.

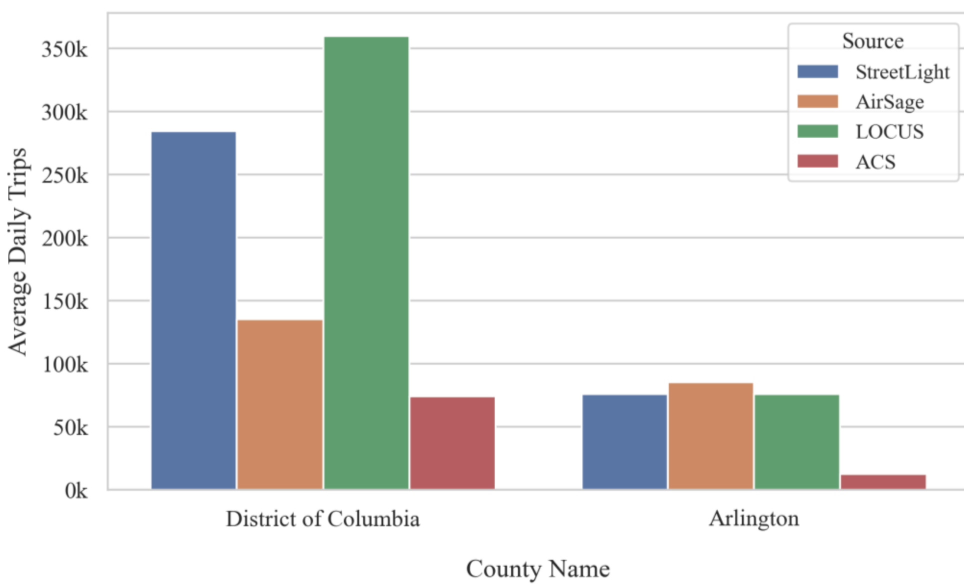

Figure 10 – Comparison of average daily pedestrian trips by county

ACS reports the lowest total volume, as expected given its commute-only definition. More importantly, ACS also shows a strong concentration of walking commute origins in the District of Columbia relative to Arlington, which is more similar to the StreetLight and LOCUS jurisdiction-level pattern than to AirSage.

The spatial maps provide additional context for the jurisdiction-level differences discussed above by illustrating where pedestrian trip origins are concentrated within each dataset. Figure 11 uses the same choropleth visualization scheme introduced previously for bicycle datasets based on quintiles. These maps compare relative spatial patterns across datasets while minimizing the influence of differences in total trip magnitude.

The percentile-based maps show broad agreement among the vendor datasets in central Washington, DC. StreetLight, AirSage, and LOCUS all identify a similar core of high pedestrian activity, consistent with the concentration of employment, commercial activity, transit access, and pedestrian-oriented destinations in that area. However, the relative patterns become more variable outside the central DC core. AirSage highlights several tracts in Arlington more prominently than StreetLight and LOCUS, consistent with the county-level results showing a more balanced distribution of pedestrian activity between the District of Columbia and Arlington. In these secondary activity areas, the vendors identify different tracts as relatively high activity, suggesting greater methodological differences in how pedestrian activity is distributed outside the highest-density urban centers.

ACS shows a more limited and less consistent spatial pattern than the vendor datasets. This is expected because ACS represents walking commute origins rather than total pedestrian activity and does not report reliable walking commute estimates for all tracts. As with the bicycle analysis, ACS is therefore useful as a contextual comparison, but it should not be interpreted as a comprehensive reference dataset for total pedestrian trips.

Population normalization was also examined for the pedestrian datasets to determine whether differences in Census tract population were influencing the mapped patterns. Dividing average daily trip origins by tract population produced only modest changes in the spatial rankings. Central Washington, DC remained the principal concentration of pedestrian activity across the vendor datasets, while differences among StreetLight, AirSage, and LOCUS outside the urban core were still evident. Some low-population tracts moved into higher per-capita classes, and tracts with no reported residential population were omitted because a rate could not be calculated. Overall, the normalized analysis reinforced rather than changed the interpretation of the trip-volume maps and is therefore not shown separately.

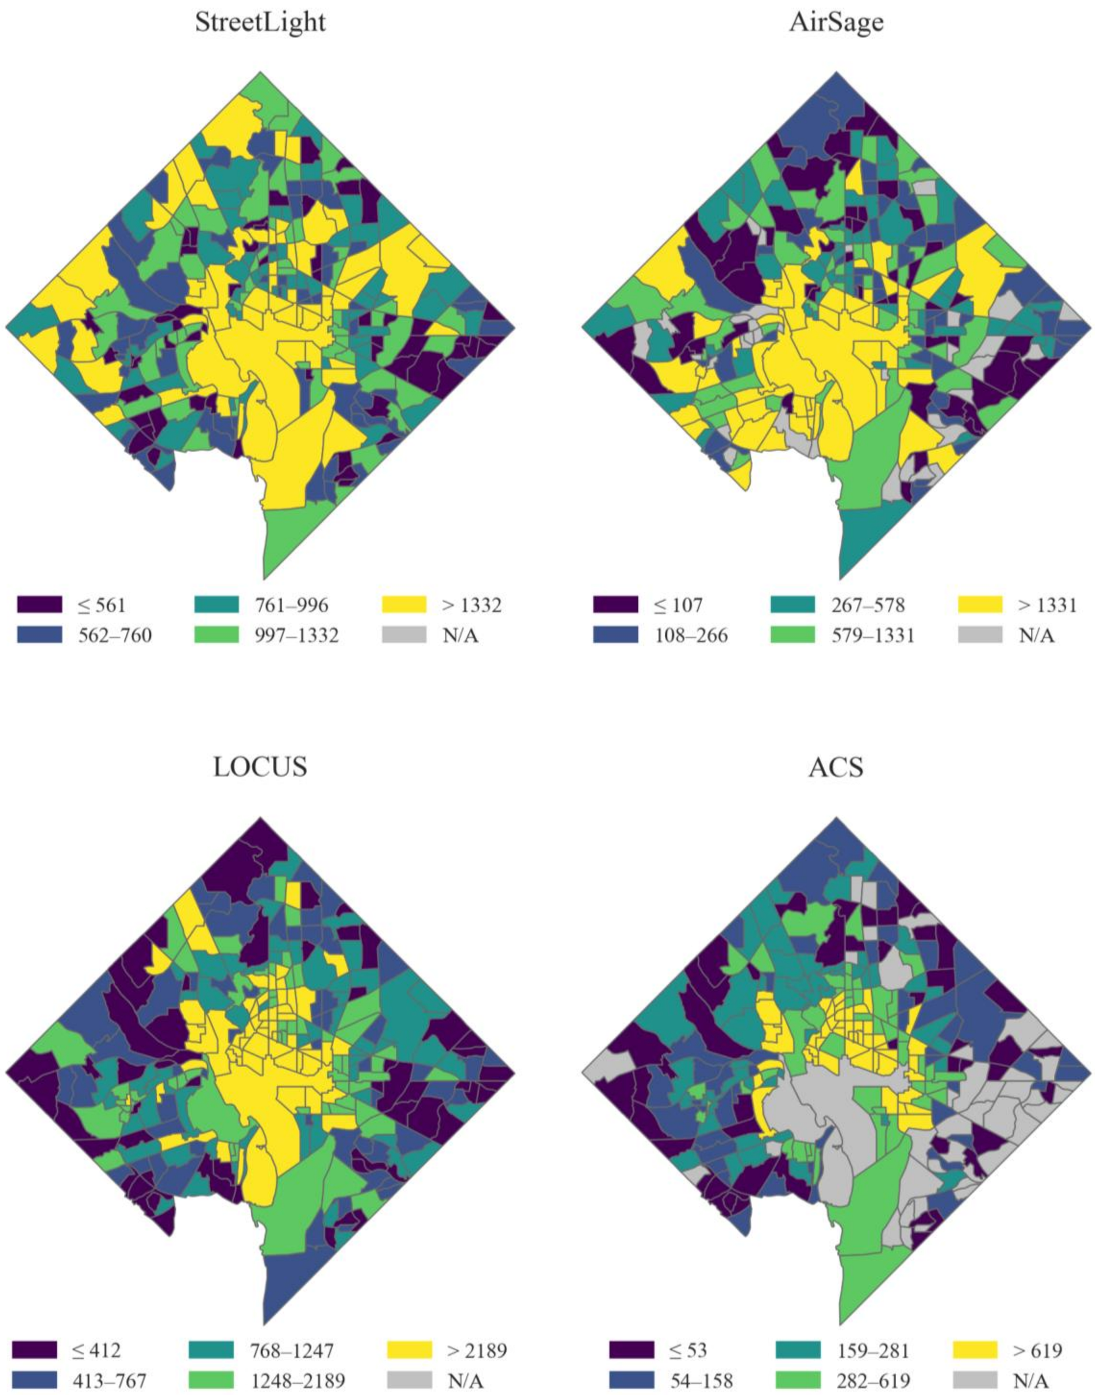

Figure 11 – Relative pedestrian activity by census tract using dataset-specific percentile classes

Figure 12 highlights the top 10 percent of Census Tracts based on average daily pedestrian trips reported by each data source. All three vendor datasets identify central Washington, DC as a major pedestrian activity area, but they differ in how concentrated or dispersed their top-ranked tracts are. LOCUS is the most concentrated around the central core, while StreetLight includes several additional high-activity tracts farther from the core DC areas. AirSage also identifies the central DC core but highlights some distinct tracts in Arlington that are not selected by StreetLight or LOCUS. These differences are consistent with the earlier spatial maps and suggest that vendor agreement is strongest for the primary pedestrian activity core, while secondary high-activity areas vary more across products.

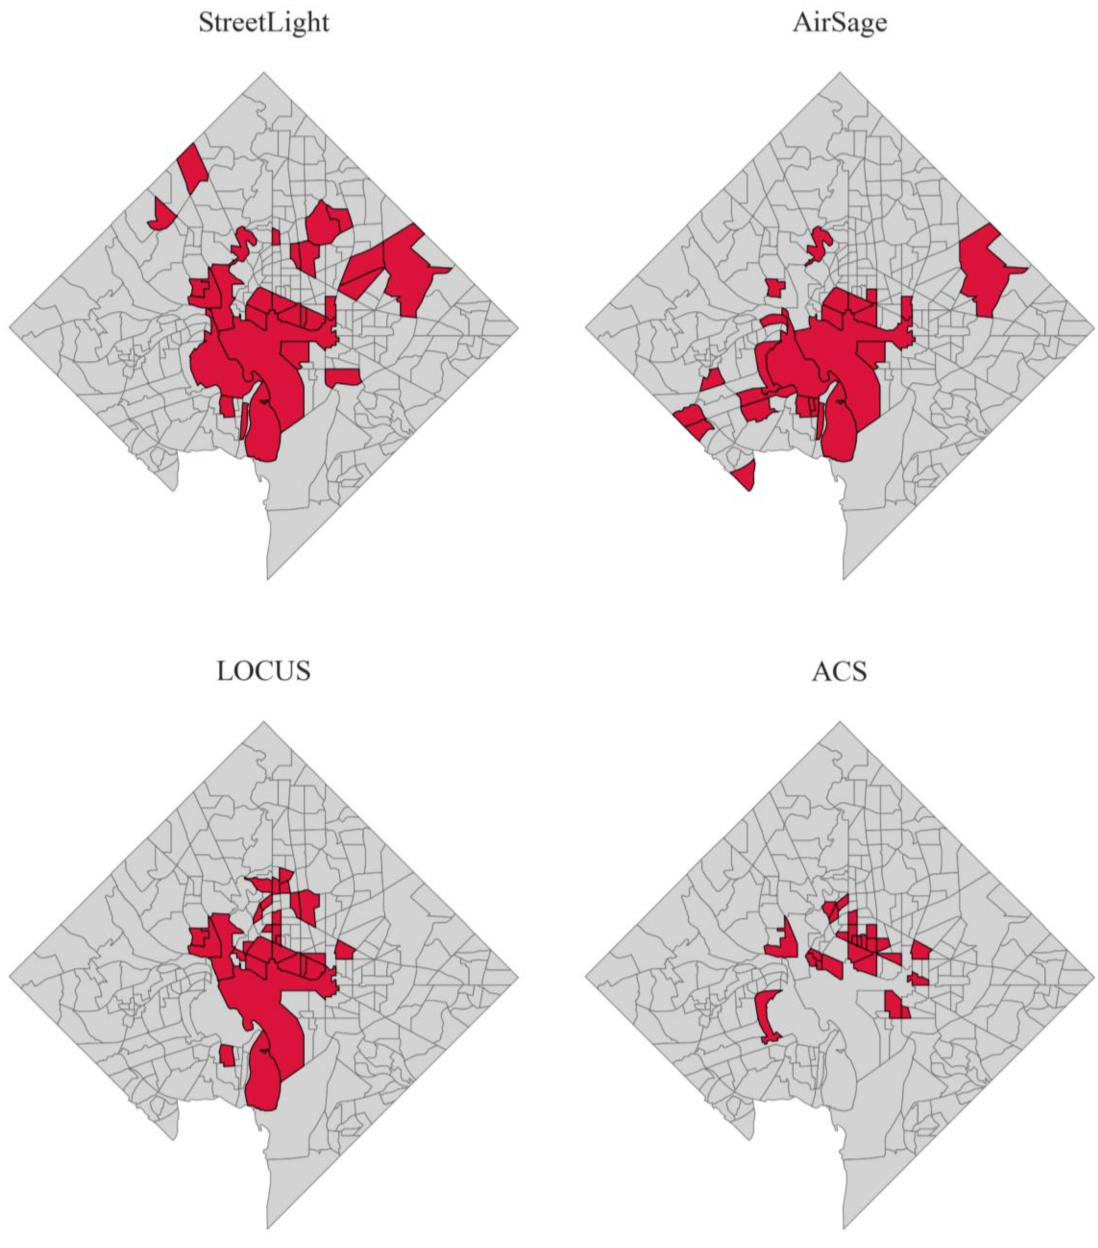

Figure 12 – Top 10% census tracts by average daily pedestrian trips

Figure 13 presents Spearman rank correlation coefficients across all census tracts for the four pedestrian datasets, providing a summary measure of spatial agreement that complements the top-10-percent tract overlap analysis above. StreetLight and LOCUS exhibit the strongest correlation (0.70) among vendors, indicating that these two datasets rank tract-level pedestrian activity more similarly than with AirSage. AirSage shows moderate correlation with StreetLight (0.50) and weaker correlation with LOCUS (0.39), consistent with the spatial maps showing that AirSage identifies a partially distinct set of high-activity tracts, particularly in Arlington County.

ACS exhibits a stronger correlation with LOCUS (0.72) than with StreetLight (0.36) or AirSage (0.22). This pattern suggests that ACS walking commute origins align more closely with the spatial ranking produced by LOCUS than with the other two datasets at the tract level. However, because ACS captures only walking commute origins rather than total pedestrian activity, correlations should not be interpreted as indicative of overall agreement in pedestrian trip volumes.

Overall, the zone-level results show that the vendor datasets generally identify central Washington, DC as a major concentration of pedestrian activity, but they differ in total reported volume and in the spatial distribution of activity outside the urban core. These comparisons should also be interpreted in light of differences in vendor product definitions, reporting periods, and day-type aggregation as described in Section 4.2. The results are most useful for understanding broad spatial consistency, relative differences across products, and areas of vendor divergence.

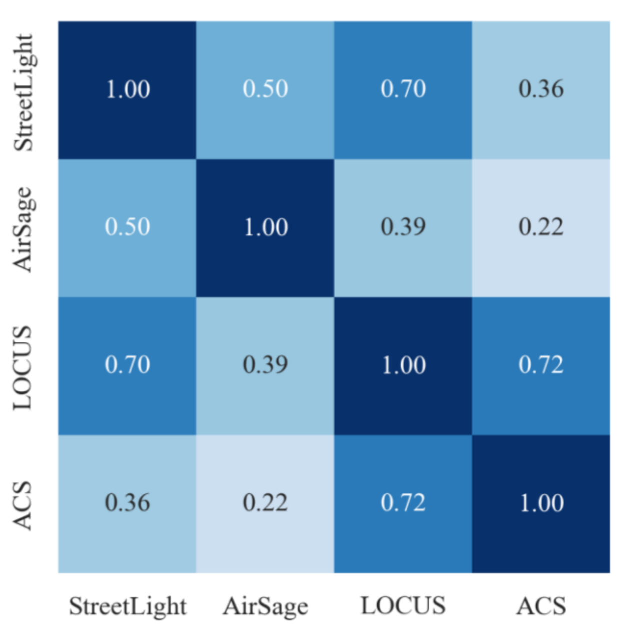

Figure 13 – Spearman correlation of tract-level average daily pedestrian trips across data sources

The rank-correlation analysis was repeated using pedestrian trip origins normalized by Census tract population. The resulting coefficients were close to those based on total trip volumes and led to the same general interpretation. StreetLight and LOCUS continued to show the strongest agreement among the vendor datasets, with a correlation of 0.69, while AirSage remained less closely aligned with both StreetLight and LOCUS. The ACS-LOCUS correlation increased slightly from 0.72 to 0.74, and the ACS-StreetLight correlation increased from 0.36 to 0.43, whereas the ACS-AirSage correlation remained unchanged at 0.22. Overall, accounting for tract population did not materially alter the relative relationships among the datasets, so the normalized correlation matrix is not presented separately.

Analysis #1 (Zone-Level) Summary

- **Bike Datasets:** Vendor datasets showed similar relative spatial patterns to one another and to the Capital Bikeshare comparative dataset, although the magnitude of trip counts differed across vendors significantly.
- **Pedestrian Datasets**: Vendor datasets generally identified central Washington, DC as a major concentration of pedestrian activity, but differed in total reported volume and in the spatial distribution of activity outside the urban core.

These results indicate that zone-level data is likely most suitable for applications that require relative measures of bike/ped activity, rather than absolute trip volumes. Bike datasets are more consistent than pedestrian datasets (across vendors), but minimally, both appear useful for identifying the highest areas of activity.

Analysis #2: Segment-Level Comparison (Network-Wide)

The network-wide segment-level analysis compares bicycle activity reported along transportation network segments. This analysis includes StreetLight and Replica as the two vendor datasets that provide segment-level bicycle volume estimates, with Strava Metro and Ride Report included as comparative datasets. StreetLight and Replica report population-level estimates of bicycle activity, while Strava Metro and Ride Report reflect observed sample counts from specific user populations. As a result, the comparative datasets are used to provide additional context for spatial patterns, rather than as ground truth measures of total bicycle activity. Note that although Ride Report is technically an INRIX product, it is used as a point of comparison here, representing a subset of shared-mobility trips (Capital Bikeshare e-bikes only).

As described in Section 4.3, direct segment-to-segment comparison is difficult because each dataset uses its own segment network and reporting structure. StreetLight reports bicycle activity on simplified bidirectional segments, while Replica uses a more detailed directional network. Strava Metro and Ride Report also rely on their own segment definitions. In addition, the datasets differ slightly in time period and day-type aggregation as described in Section 4.2. To make the datasets as comparable as possible, segment-level measures were standardized to average daily trips, and directional estimates were converted to bidirectional values where needed. However, these standardization steps do not fully resolve differences in network representation, reporting assumptions, or the underlying populations captured by each dataset. Therefore, the analysis does not treat individual segments as directly equivalent across datasets. Instead, the network-wide comparison focuses on two broad forms of analysis.

First, distributional summaries are used to compare how bicycle activity is distributed across segments within each dataset. Second, maps of high-activity segments are used to visually assess whether the datasets identify similar corridors and locations with concentrated bicycle activity. These comparisons provide a broad indication of spatial consistency across products while recognizing the limitations created by different network definitions, reporting assumptions, and the absence of network-wide ground truth.

In addition to these bicycle datasets, the analysis also examines INRIX's Vulnerable Road User (VRU) Index. As described in Section 4.2, the VRU Index differs fundamentally from the other segment-level products. Rather than estimating bicycle volumes, it reports the relative concentration of combined bicycle and pedestrian activity on a roadway segment compared with other segments in the study area. The metric reflects relative activity rather than trip counts, combines multiple active transportation modes into a single measure, and excludes off-road trails and paths because it is reported only on roadway segments that permit vehicle traffic.

Consequently, the VRU Index is not directly comparable to the bicycle volume datasets and is evaluated separately as a measure of relative active transportation activity. INRIX reports directional VRU Index values on a highly granular roadway network. For visualization purposes and to improve consistency with the other datasets, directional VRU Index values were converted to bidirectional values by averaging the index reported for each direction of travel along the same roadway segment.

Comparing Trip Distributions

Table 4 and Figure 14 compare the distribution of segment-level average daily bicycle trips across StreetLight, Replica, Strava Metro, and Ride Report. These comparisons are intended to describe how bicycle activity is distributed within each dataset, rather than to directly compare one dataset against another. This distinction is important because in addition to the differences mentioned, the number of reported segments differs substantially across datasets, reflecting differences in network coverage, segmentation, and facility representation.

Table 4 – Summary statistics for segment-level average daily bicycle trips by data source

Vendor	N Segments	Network Length (mi)	Mean	SD	Min	25th Percentile	50th Percentile (Median)	75th Percentile	Max
StreetLight	22,957	760.8	95.9	210.3	0	0	16	102	2,912
Replica	45,975	874.8	170.9	455.8	1	4	15	83	4,950
Strava	61,952	1050.2	4.1	14.3	0	0	1	3	477
Ride Report	22,416	655.2	16.2	25.0	0	0	5	22	224

N Segments: Number of segments; SD: Standard deviation; Network Length: Represents the total length of segments included in each vendor's submission within the study area

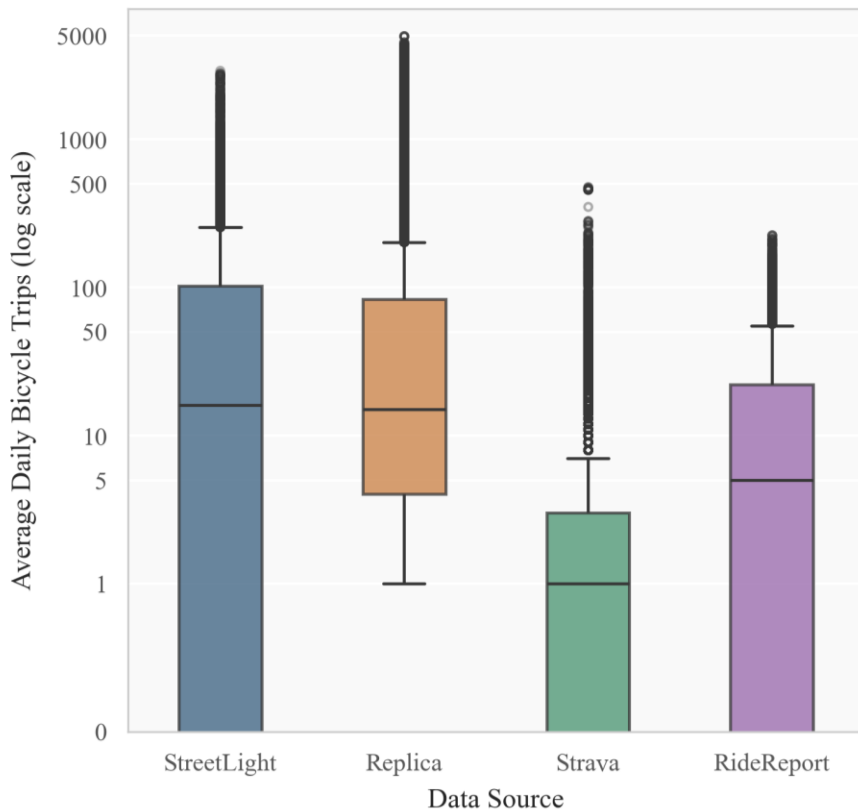

Figure 14 – Distribution of segment-level average daily bicycle trips by data source

StreetLight reports 22,957 segments, with a mean of 95.9 average daily bicycle trips and a median of 16. The large difference between the mean and median indicates a strongly right-skewed distribution, with many low-volume segments and a smaller number of segments carrying substantially higher bicycle activity. This pattern is also reflected in the distribution percentiles: 25 percent of StreetLight segments report zero bicycle trips, while the 75th percentile reaches 102 trips per day. The maximum reported value is 2,912 trips per day, indicating the presence of a relatively small number of high-volume corridors.

Replica reports the highest mean segment volume, approximately 171 average daily bicycle trips, but has a median of only 15 trips per day. Similar to StreetLight, this large difference between the mean and median indicates a highly skewed distribution. Most Replica segments have relatively low reported bicycle volumes, while a smaller number of segments carry very high estimated volumes. This is reflected in the 75th percentile of 83 trips and the maximum value of 4,950 trips. The boxplot reinforces this pattern, showing relatively low central distribution but a large number of high-volume outliers. Compared with Replica, StreetLight exhibits a somewhat higher concentration of moderate-volume segments, as indicated by its higher 75th percentile, but both datasets display broadly similar right-skewed distributions.

The two comparative datasets have much lower segment-level values, which is expected because they represent sample counts rather than population estimates. Strava Metro includes the largest number of segments, 61,952, but has a mean of only 4.2 average daily bicycle trips and a median of 1. The 25th percentile is 0 and the 75th percentile is 3, indicating that most Strava segments have very low observed activity during the analysis period. Ride Report shows somewhat higher values than Strava, with a mean of 16.2 and a median of 5, but it also has a 25th percentile of 0.

This indicates that a substantial portion of the Ride Report network has no or very limited reported activity, while a smaller set of segments has more frequent observations.

The distributional results suggest that StreetLight and Replica allocate bicycle activity across their networks in broadly similar ways, with both datasets characterized by many low-volume segments and a smaller number of high-volume corridors. Differences remain in the extent to which activity is concentrated on the highest-volume segments, which may reflect a combination of network representation, underlying data sources, and estimation methodologies. Strava Metro and Ride Report provide useful comparative context for observed bicycle activity patterns, but their lower values reflect their role as sample-based datasets and should not be interpreted as directly comparable to the vendor population estimates. The vast difference in absolute volumes reported among the different data sources points to the need to compare highest activity segments, as shown in the next section.

Comparing High Activity Corridors

The high-activity segment maps provide a visual comparison of where each dataset places its highest bicycle activity segments. Figure 15 through Figure 17 show the top 1 percent, 5 percent, and 10 percent of segments, respectively, based on average daily bicycle trips within each dataset. The maps focus on the core portions of Washington, DC and Arlington, Virginia to improve readability and allow corridor-level patterns to be compared more clearly. Because the volume thresholds are defined separately within each dataset, these maps should be interpreted as relative comparisons of each dataset's highest-activity segments, rather than direct comparisons of absolute bicycle volume.

Across the four datasets, the highlighted segments generally form recognizable corridors rather than appearing as isolated or randomly scattered segments. This is an important qualitative finding because it suggests that, within each dataset, higher bicycle activity tends to be assigned in a spatially coherent way along connected facilities and travel corridors. The degree and location of this continuity differ across datasets, but the overall presence of corridor-like patterns supports the use of these maps for broad visual comparison.

- **StreetLight** shows a distinct pattern in which many of the highest-ranked segments are concentrated in and around the National Mall, central Washington, DC, and major connections near the Potomac River. In the top 1 percent map, the highlighted segments are concentrated along the National Mall, riverfront corridors, and several major bridge and trail connections. As the threshold expands to the top 5 percent and top 10 percent, StreetLight identifies a broader set of corridors throughout the urban core, including additional east-west and north-south connections. Nevertheless, the strongest concentration of high-activity segments remains centered on the National Mall, downtown Washington, DC, and Potomac River crossings. Compared with the other datasets, StreetLight places greater emphasis on these central recreational and commuter corridors.

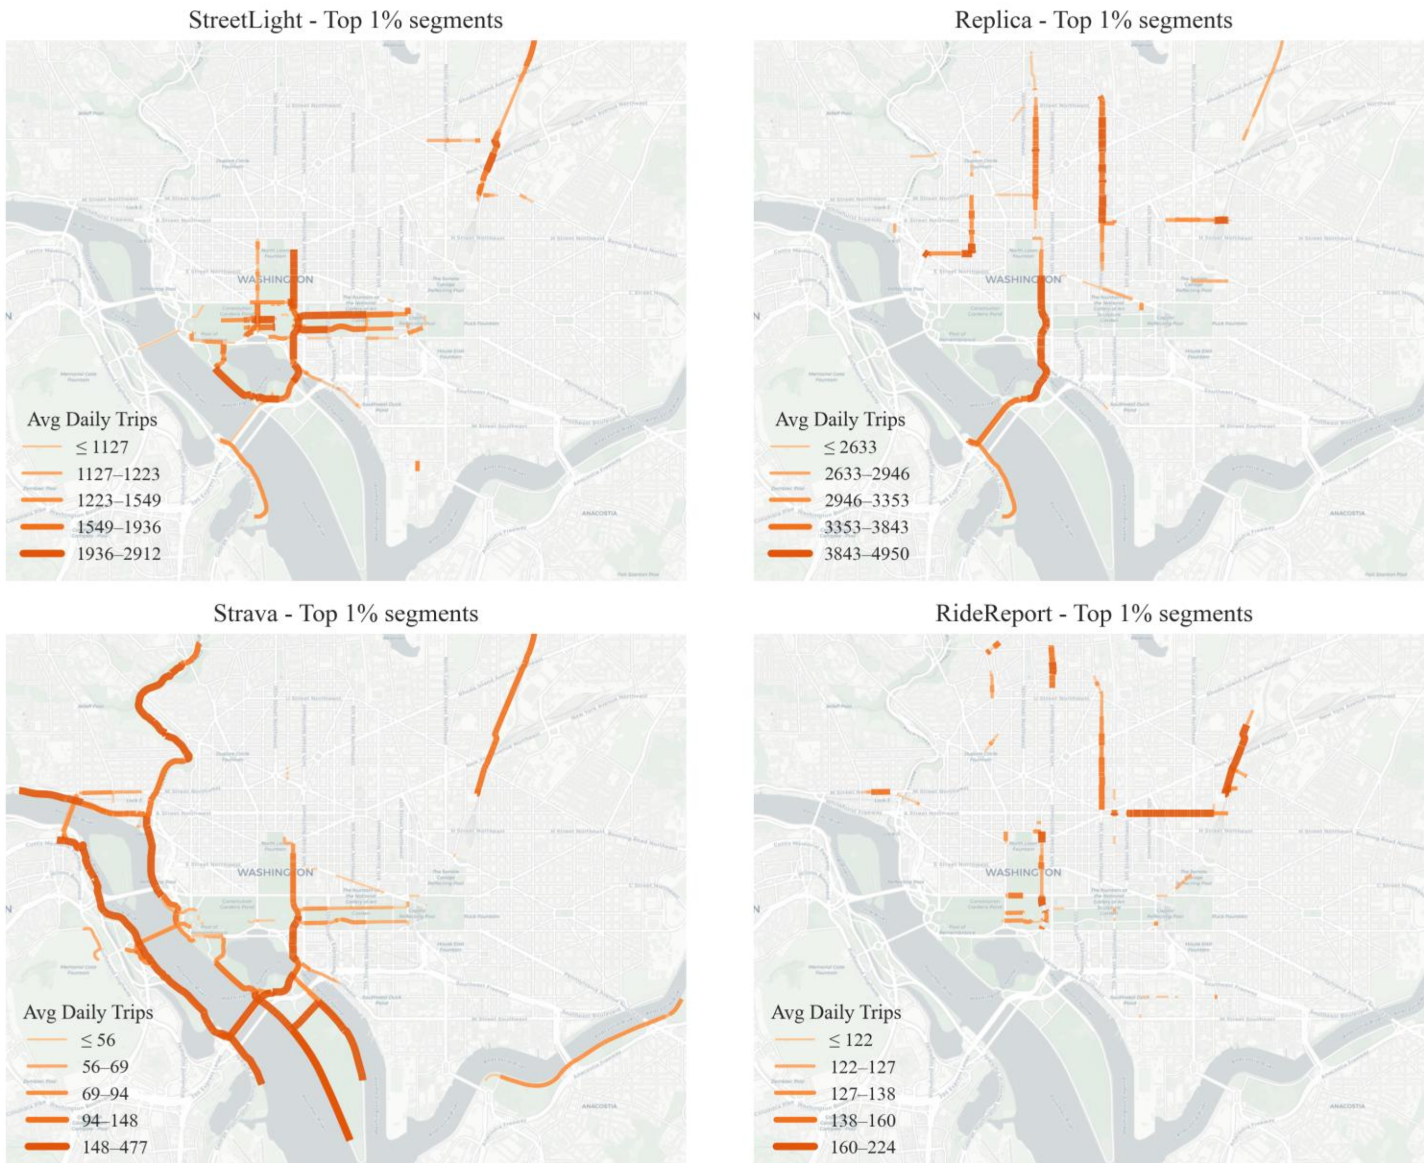

Figure 15 – Top 1 percent of segment-level average daily bicycle trips by data source

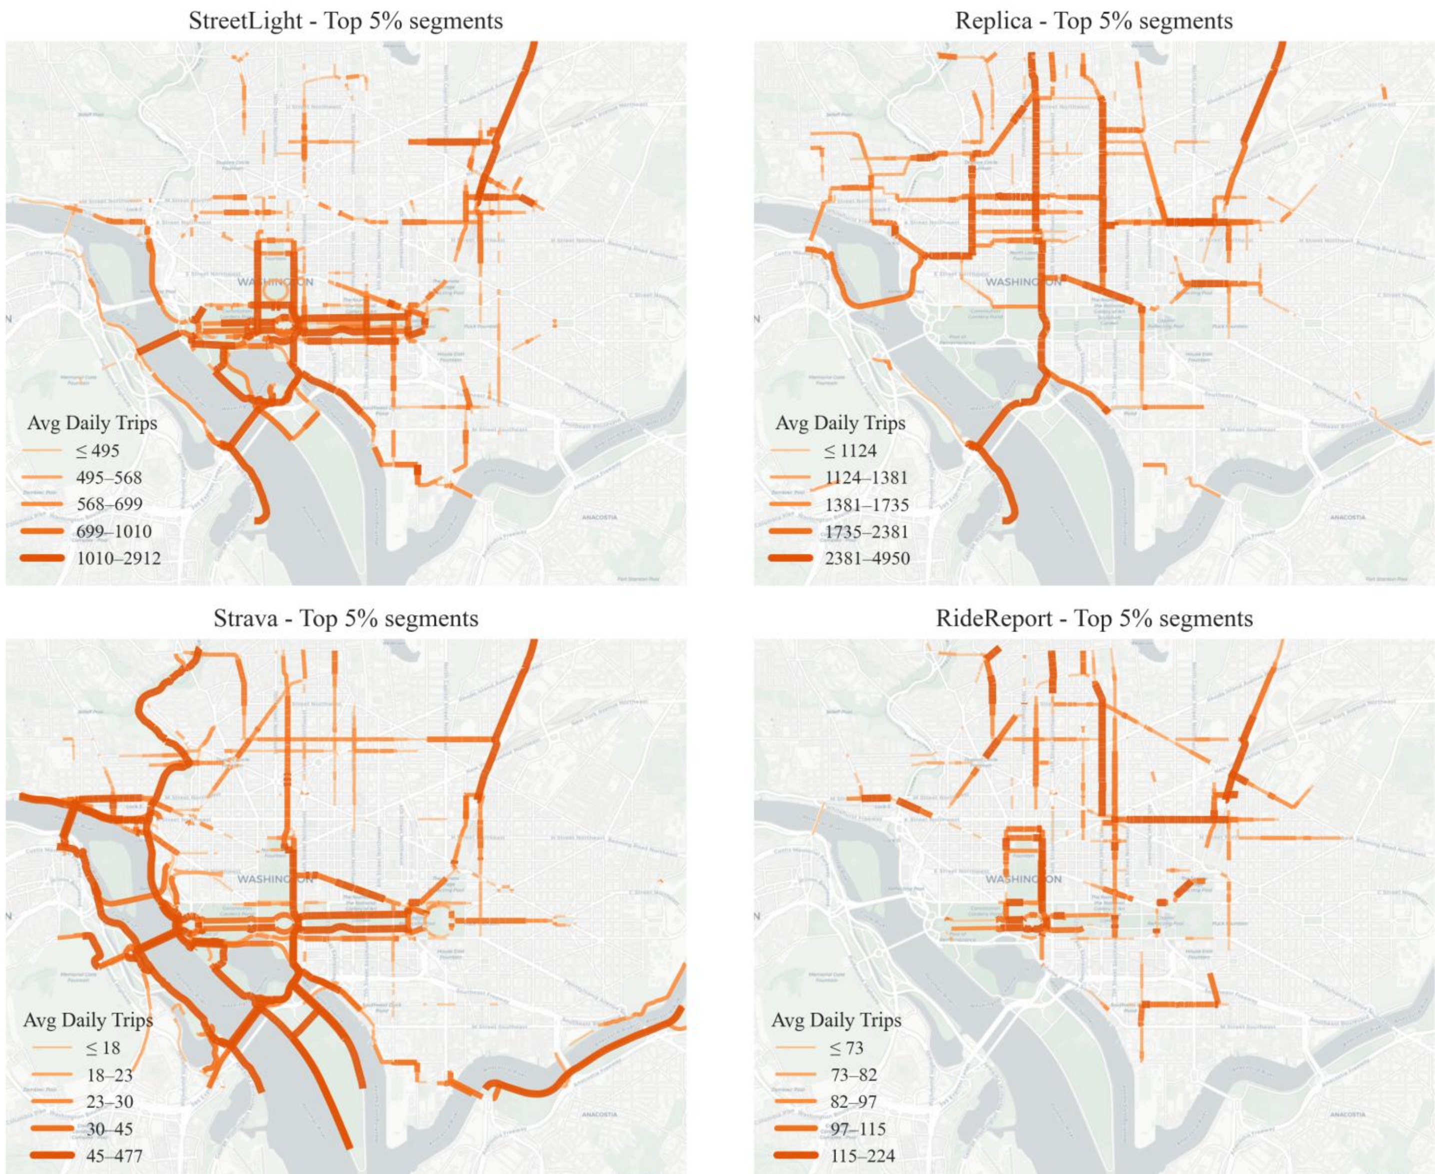

Figure 16 – Top 5 percent of segment-level average daily bicycle trips by data source

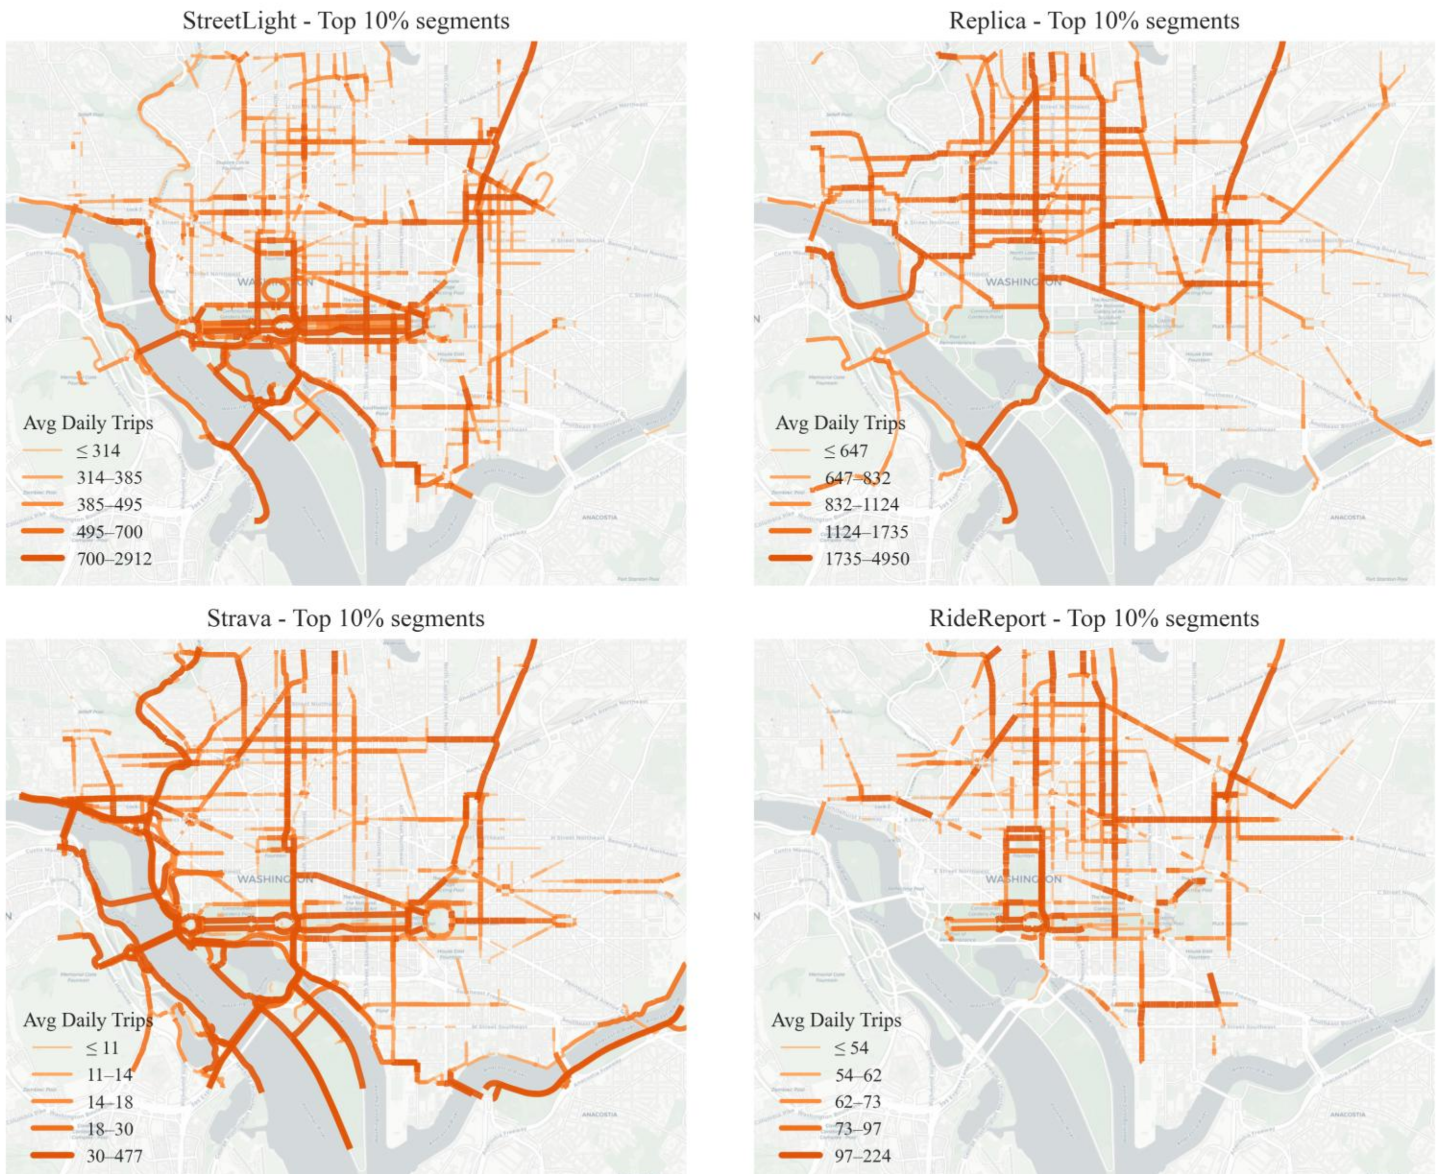

Figure 17 – Top 10 percent of segment-level average daily bicycle trips by data source

- **Replica** presents a notably different pattern. Its top segments are less concentrated around the National Mall and instead form a broader grid-like distribution across the study area. This pattern is visible even in the top 1 percent map and becomes more pronounced in the top 5 percent and top 10 percent maps, where high-activity segments extend across a larger portion of the downtown street grid and surrounding neighborhoods. Relative to StreetLight, Replica distributes high bicycle activity across a wider set of urban corridors rather than concentrating activity on a smaller number of prominent regional facilities.
- **Strava Metro** shows a strong emphasis on trail and recreational cycling corridors. High-activity Strava segments are prominent along major trails and river-adjacent routes, including corridors around the Potomac River and connections into central Washington, DC. This pattern is consistent with Strava's user base, which is more likely to include recreational and fitness-oriented cyclists. At the same time, Strava also identifies some central urban bicycle corridors, including areas around the National Mall and selected north-south facilities that also appear in other datasets, such as corridors highlighted by Replica. This suggests that Strava provides useful comparative evidence for major bicycle routes, particularly trails and well-known cycling corridors, but may place greater relative emphasis on recreational facilities than the vendor population estimates.
- **Ride Report** highlights a set of high-activity corridors that are more concentrated within the urban street network than along major trail facilities. This pattern is consistent with its data source, which is more closely tied to shared micromobility activity, including e-scooter and e-bike trips. Several Ride Report corridors appear visually similar to those highlighted by Replica, particularly within the downtown street grid, and Ride Report also identifies multiple segments in or near the National Mall. At the same time, major trail corridors are less prominent in Ride Report than in some of the other comparative patterns, reinforcing that it should be interpreted as a complementary source reflecting a specific segment of bicycle and micromobility activity rather than a comprehensive measure of all bicycle travel.

The high-activity segment maps show meaningful differences across the datasets, particularly between the two vendor products. Both StreetLight and Replica identify coherent high-activity bicycle corridors, and both consistently highlight major facilities near the National Mall and Potomac River crossings. However, StreetLight places greater emphasis on these central regional corridors, while Replica distributes high bicycle activity more broadly across the urban street network. Because there is no comprehensive network-wide ground-truth dataset for bicycle volumes, the analysis cannot determine which representation is more accurate. Instead, the differences should be interpreted as evidence that the two products may produce different views of how bicycle activity is distributed across the transportation network.

As an additional exploratory check, StreetLight and Replica segments sharing a common OpenStreetMap identifier were separated into two broad facility groups: paths and cycleways, and other segments representing primarily the urban roadway network. Although this matching approach does not fully resolve differences in how the vendors represent parallel roadways and bicycle facilities, the comparison showed closer agreement between StreetLight and Replica on several major cycleways and shared-use paths than was apparent from the full-network maps. The roadway results continued to reflect the broader contrast identified above, with StreetLight emphasizing a smaller set of prominent corridors and Replica distributing high bicycle activity across a wider portion of the urban street grid. Because the common-network filter can exclude or retain different components of parallel roadway and cycleway configurations, these findings are treated as a qualitative sensitivity check and the corresponding maps are not presented separately.

INRIX VRU

Figure 18 presents the highest-ranked roadway segments according to the INRIX Vulnerable Road User (VRU) Index. Unlike the previous datasets, the VRU Index does not represent bicycle volumes. Instead, it measures the relative concentration of combined bicycle and pedestrian activity on roadway segments compared with other segments in the study area. As a result, the figure should be interpreted as a relative indicator of active transportation activity rather than a bicycle volume map. The VRU Index highlights a spatial pattern that differs from the bicycle-focused datasets. The highest-ranked segments are concentrated within the downtown Washington, DC street grid and surrounding activity centers, with less emphasis on the National Mall, Potomac River crossings, and major trail corridors. As the threshold expands from the top 1 percent to the top 20 percent of segments, the highlighted network grows outward but remains focused on dense urban roadway corridors. This pattern is more similar to Replica's grid-oriented distribution than to the corridor-focused patterns observed in StreetLight and Strava. These differences are consistent with the design of the VRU Index, which combines bicycle and pedestrian activity and is reported only on roadway segments.

Taken together, the bicycle datasets and the INRIX VRU Index provide complementary perspectives on active transportation activity. While the bicycle datasets emphasize where bicycle travel is concentrated, the VRU Index highlights roadway corridors with elevated overall vulnerable road user activity. The differing patterns suggest that the choice of dataset can influence which facilities emerge as priorities for planning, analysis, and investment

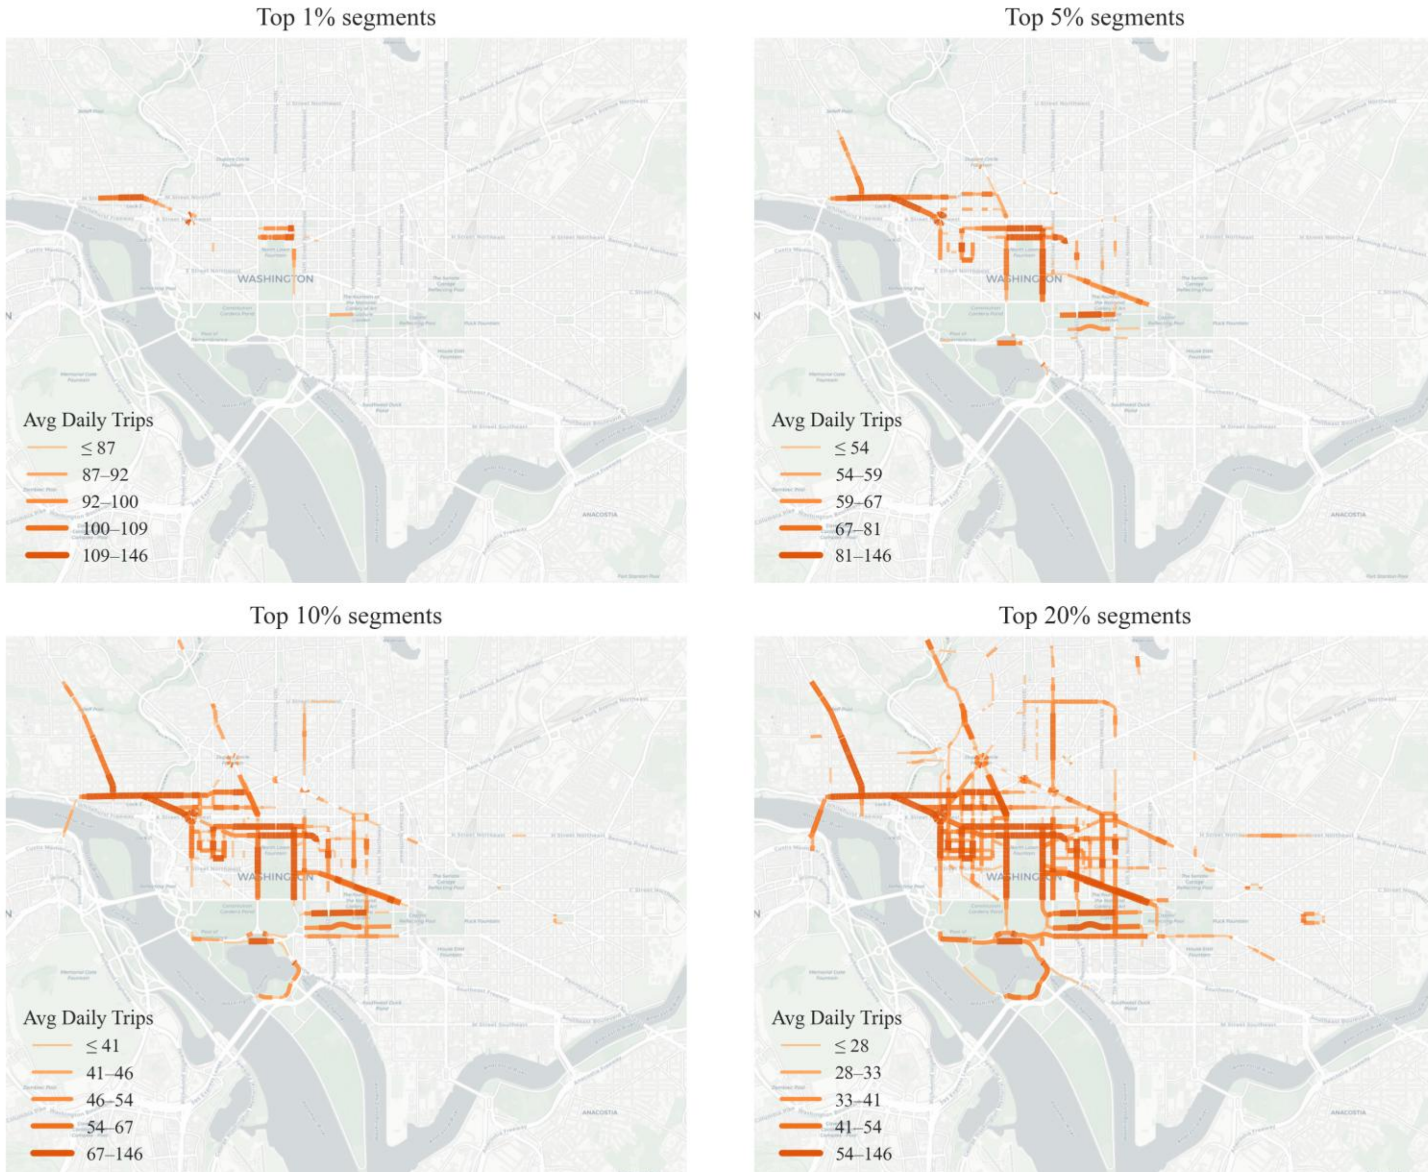

Figure 18 – High-activity roadway segments identified by the INRIX VRU Index (top 1%, 5%, 10%, and 20% of segments)

StreetLight Data Resubmission

After the initial analysis was completed, StreetLight provided an updated version of its segment-level bicycle activity dataset, coinciding with its official product release (the initial submission was a pre-release version shared for inclusion in the validation activity). Because the updated delivery was received prior to finalizing this report, all StreetLight results presented throughout the report are based on the updated dataset. This subsection is included to show the differences between the original and updated deliveries and to provide context regarding how the revised product affected the network-wide bicycle activity patterns identified by StreetLight.

Table 5 compares summary statistics for the original and updated StreetLight deliveries. The updated delivery contains fewer reported segments (22,957 versus 33,292), largely reflecting additional network simplification, particularly near intersections where multiple shorter segments in the original delivery were consolidated into longer segments in the updated network. The updated delivery also reports lower bicycle activity overall, with the mean decreasing from 123.4 to 95.9 average daily bicycle trips and the median decreasing from 75 to 16 trips. The distribution of bicycle activity changed substantially as well. In the original delivery, the 25th percentile segment reported 45 trips per day, indicating that most segments were assigned moderate bicycle volumes. In the updated delivery, the 25th percentile decreased to zero and the median fell considerably, suggesting a much larger share of low-volume segments. At the same time, the maximum reported segment volume increased from 1,556 to 2,912 trips per day, indicating that bicycle activity became more concentrated on a smaller subset of high-volume corridors.

Table 5 – Summary statistics for the original and updated StreetLight deliveries

StreetLight	N Segments	Mean	SD	Min	25th Percentile	50th Percentile (Median)	75th Percentile	Max
Original	33,292	123.4	140.2	5	45	75	137	1,556
Updated	22,957	95.9	210.3	0	0	16	102	2,912

Figure 19 compares the top 1 percent, 5 percent, and 10 percent of bicycle activity segments identified by the two deliveries. Despite the changes in segment-level distributions, the overall spatial patterns remain broadly consistent. Both deliveries identify the National Mall, Potomac River crossings, and major bicycle corridors in central Washington, DC as locations of elevated bicycle activity. However, the updated delivery places greater emphasis on several regional corridors and bridge approaches while also expanding the network of high-activity segments identified within the downtown street network. These differences are particularly visible in the top 5 percent and top 10 percent maps, where the updated delivery highlights a more extensive set of north-south and east-west urban corridors than the original delivery.

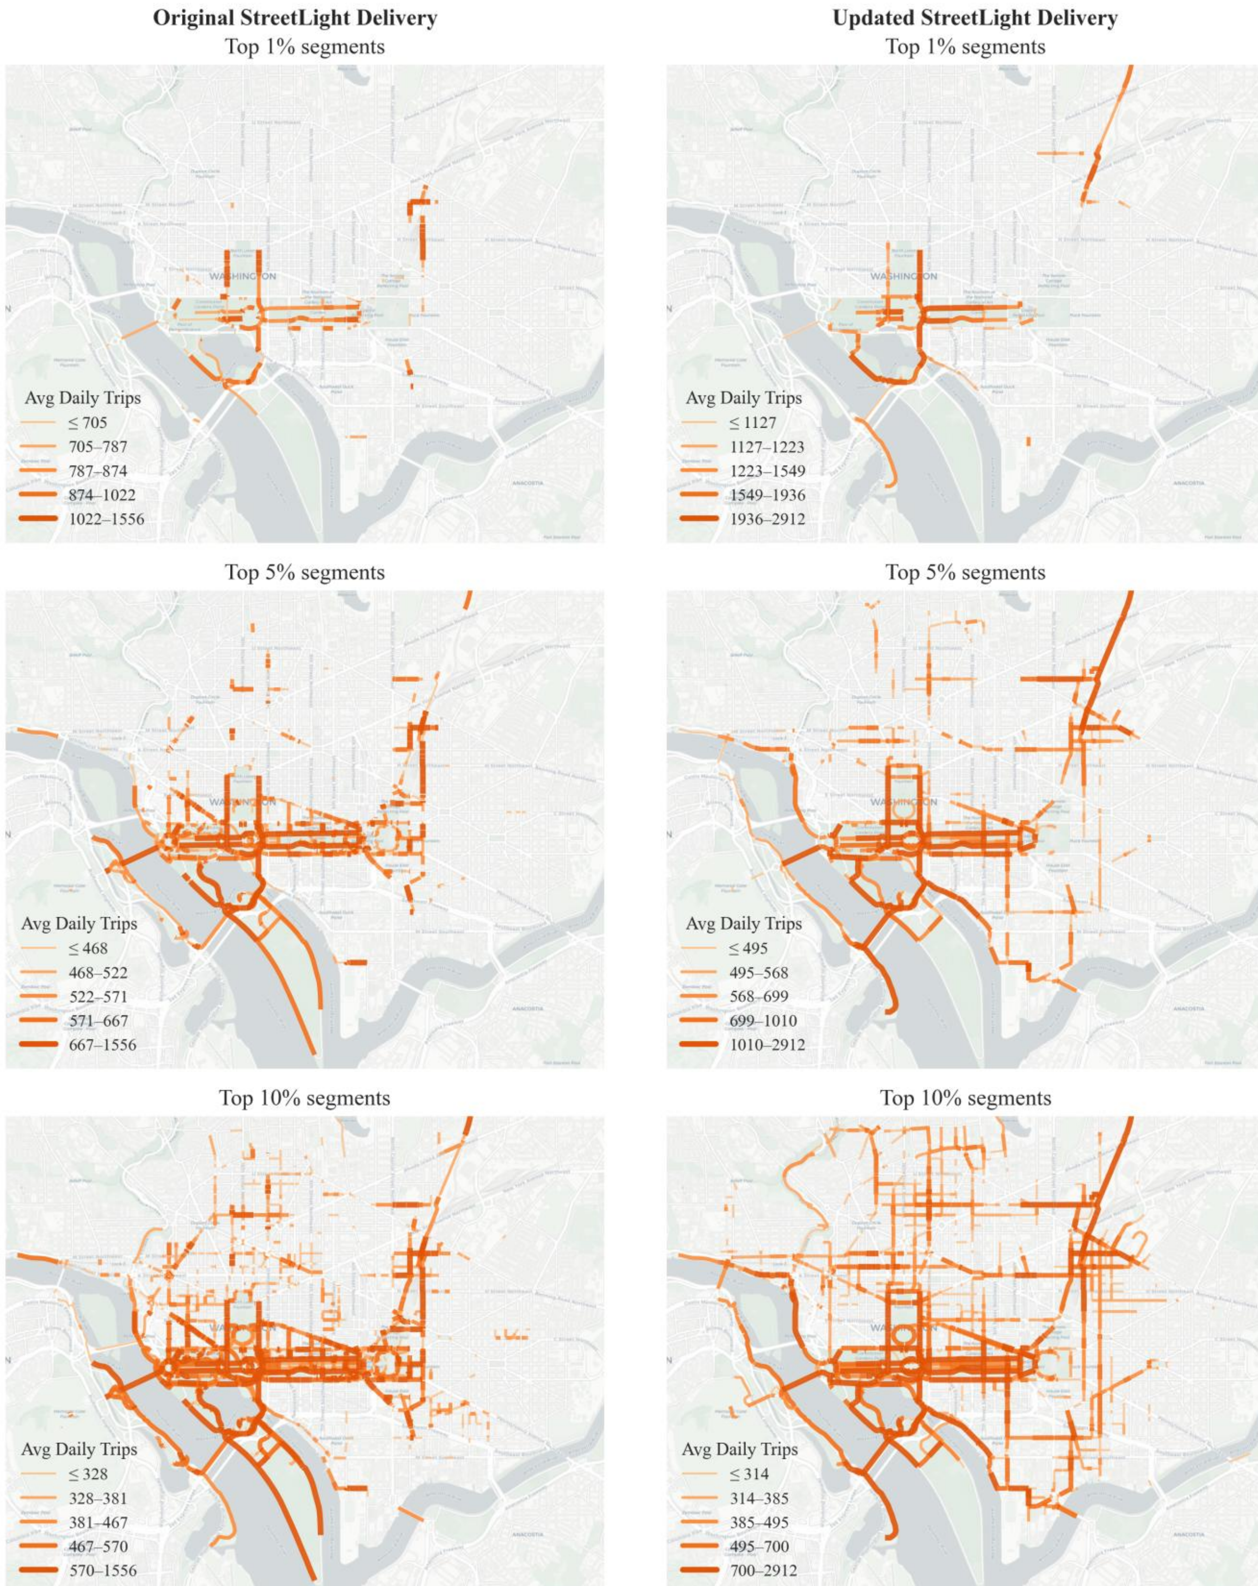

Figure 19 – Top 1%, 5%, and 10% bicycle activity segments from the original and updated StreetLight deliveries

Analysis #2 (Segment Level Comparison, Network-Wide) Summary

- **Segment Distributions:** StreetLight and Replica datasets were characterized by mostly low-volume segments, with bicycle activity concentrated on a smaller number of high-volume corridors.
- **Comparing High-Activity Corridors:** All datasets identified spatially coherent bicycle corridors, but their patterns differed. StreetLight emphasized central and regional facilities near the National Mall and Potomac River, while Replica distributed high activity more broadly across the urban street grid. Strava placed greater emphasis on trails, while Ride Report and the INRIX VRU Index highlighted more urban roadway activity.

These results indicate that segment-level products can support the identification of broad high-activity corridors, but differences in network representation, product definitions, and underlying data sources can substantially affect which facilities are prioritized.

Analysis #3: Segment-Level Evaluation at Agency Count Sites

This section evaluates vendor-reported segment-level bicycle volumes at locations where agency count data are available. The count locations used in this analysis were described previously in Section 4.2. Unlike the network-wide segment comparison, which focused on broad spatial patterns and high-activity corridors, this analysis uses observed count data as a reference point to assess how closely vendor estimates align with measured bicycle activity at specific locations.

The agency count data used in this evaluation were obtained from publicly available sources, including the BikePed Portal. Although the validation team did not provide these count data to vendors as part of this study, portions of the historical count records were publicly accessible and may have been available to vendors independently. Such data could potentially be used by vendors for product development, calibration, or evaluation purposes. However, vendor methodologies and calibration practices are proprietary, and the extent to which any specific vendor may have incorporated these data is unknown.

Preparing the data for this comparison required matching the count data to each vendor's reporting format as closely as possible. A crosswalk between agency count locations and the corresponding vendor segments was created manually through careful review of the vendor networks and count-site locations. This step was necessary because StreetLight and Replica use different segment networks, and count sites do not always align directly with vendor segment definitions. The matching process was therefore conducted separately for each vendor to identify the most appropriate segment-count pairs for comparison.

The temporal aggregation of the count data was also tailored to each vendor product. StreetLight reports annual average daily bicycle volumes for 2024, so corresponding count-site average daily volumes were calculated for the same annual period where sufficient count data were available. Because StreetLight reports bidirectional segment volumes, comparison cases were limited to locations where count-site volumes were available for both directions. The directional count volumes were then aggregated to produce a bidirectional observed value consistent with StreetLight's reporting format. This requirement reduced the number of valid StreetLight comparison cases. Replica reports directional segment-level average daily bicycle volumes for fall 2024. Therefore, for Replica, the count-site volumes were calculated for the corresponding fall 2024 period and matched to directional vendor segments. This allowed the comparison to preserve the directional structure of Replica's reported volumes, rather than aggregating all observations to a bidirectional value.

Because the StreetLight and Replica comparisons differ in temporal aggregation, network matching, and directional structure, the results are presented separately. StreetLight's results are discussed first, followed by Replica's.

StreetLight Count-Site Comparison

For StreetLight, the count-site comparison includes 33 cases where agency count volumes were available for both directions and could therefore be aggregated to match StreetLight's bidirectional segment estimates. Figure 20 compares counter-based average daily bicycle trips with StreetLight's reported values. The Spearman correlation of 0.76 indicates a strong positive relationship, suggesting that locations with higher observed bicycle activity generally receive higher StreetLight estimates. While some site-level differences remain, the overall agreement is substantially stronger than would be expected from a purely qualitative ranking measure. The mean absolute percentage error of 57 percent, mean absolute error of 192 trips per day, and RMSE of 257 trips per day indicate moderate site-level error, but overall, the dataset captures broad differences in bicycle activity across count locations reasonably well.

Figure 21 plots the difference between StreetLight and counter volumes against the observed counter volume. A LOESS (Locally Estimated Scatterplot Smoothing) trend line is included to reveal systematic patterns in the differences without imposing a predefined functional form. The mean bias is -36 trips per day, indicating that StreetLight slightly underestimates bicycle activity on average, although the magnitude of this bias is relatively small compared with the observed volume range. The LOESS trend remains relatively close to zero across most of the volume range and does not exhibit a strong systematic pattern of over- or under-estimation. This observation is supported by the proportional bias statistic ($r = -0.20$, $p = 0.273$), which indicates that the magnitude of estimation error is not significantly associated with the observed bicycle volume. In other words, there is little evidence that StreetLight's performance systematically changes across low- and high-volume count locations.

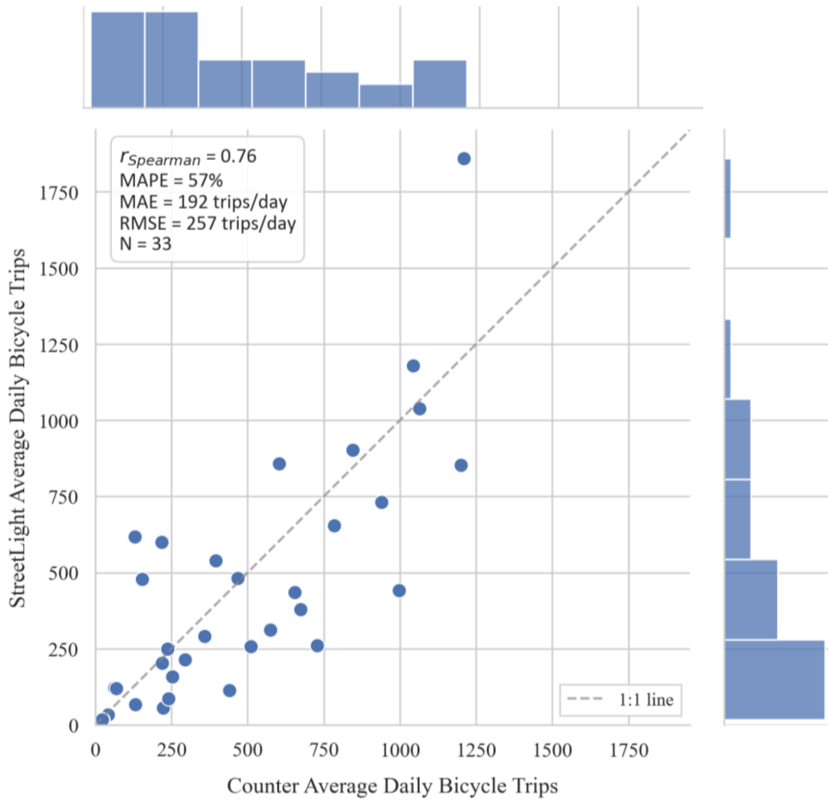

Figure 20 – Comparison of StreetLight and agency count average daily bicycle volumes

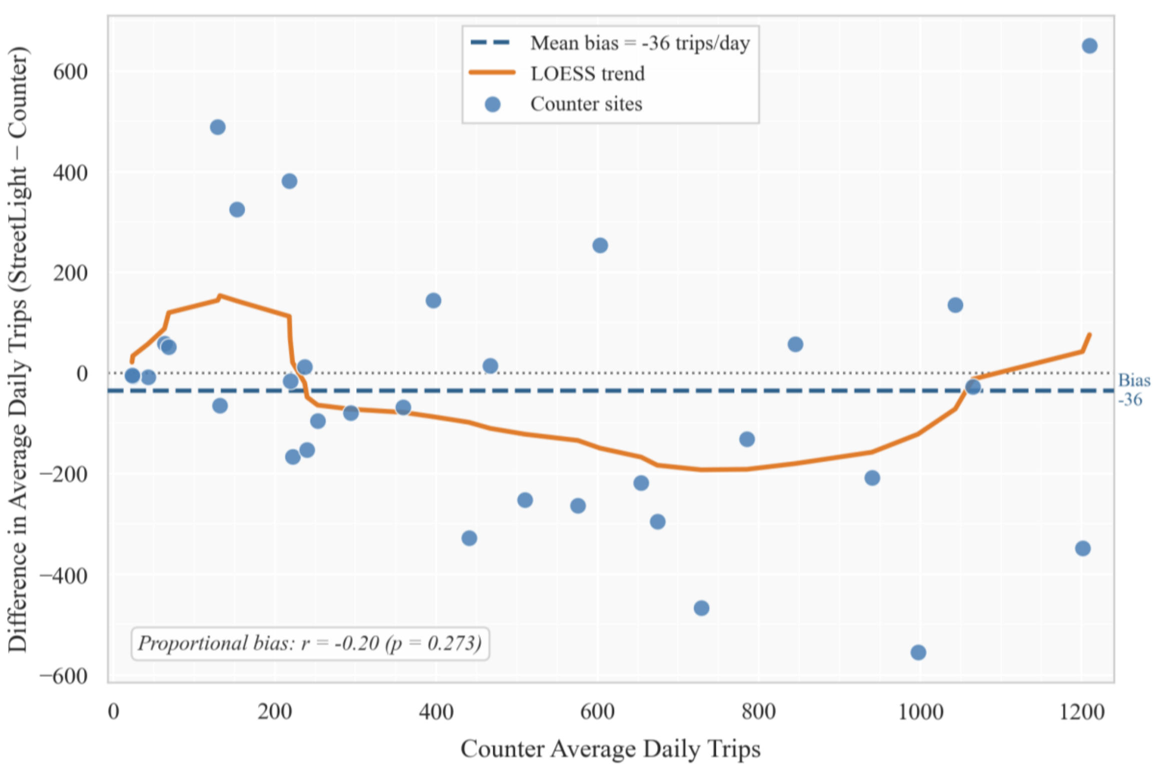

Figure 21 – Difference between StreetLight and agency count volumes by count-site volume

Table 6 and Figure 22 and Figure 23 examine how StreetLight errors vary by observed bicycle volume range and facility type. Error is defined as StreetLight minus counter volume, so negative values indicate underestimation by StreetLight. Absolute percentage error (APE) measures the size of the difference relative to the counter volume, regardless of direction.

Table 6 – StreetLight error summary by counter volume range and facility type

Scenario		N	Error (trips)				Absolute Percentage Error (%)			
			Mean	P25	P50	P75	Mean	P25	P50	P75
Overall		33	-36	-208	-27	58	57%	20%	38%	63%
By Volume Range	**<150**	7	73	-8	-5	55	95%	25%	49%	84%
	150-500	12	-2	-110	-42	47	61%	16%	37%	74%
	>500	14	-119	-287	-213	37	34%	18%	38%	49%
By Facility Type	**Road**	8	-123	-156	-46	-8	40%	21%	39%	58%
	Trail	25	-8	-218	-16	136	62%	19%	38%	64%

N: Number of cases; P25: 25th percentile; P50: 50th percentile (Median); P75: 75th percentile

The boxplots provide a visual summary of how errors are distributed within each group. In each boxplot, the box represents the middle 50 percent of observations, from the 25th percentile to the 75th percentile, also known as the interquartile range. The horizontal line inside the box represents the median, or 50th percentile. The whiskers extend beyond the box to show the estimated non-outlier range, up to 1.5 times the interquartile range below the 25th percentile and above the 75th percentile. Individual comparison cases are shown as dots overlaid on each boxplot, allowing the site-level observations within each group to be seen directly. Points beyond the whiskers represent cases that fall outside the typical range for that group. The P25, P50, and P75 values in Table 6 correspond to the same percentile values shown visually in the boxplots.

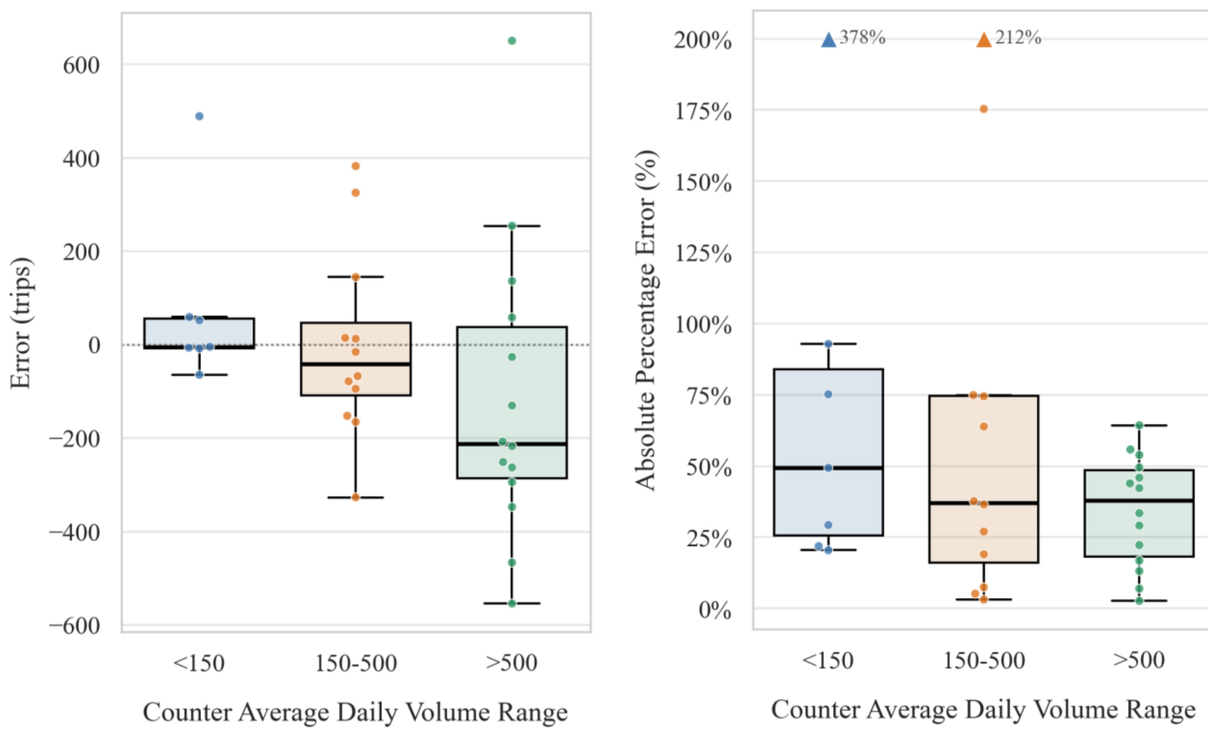

Figure 22 – Distribution of StreetLight errors and absolute percentage errors by counter volume range

The volume-range comparison reveals notable differences in both the magnitude and variability of errors across count locations. For sites with fewer than 150 average daily bicycle trips, StreetLight slightly overestimates activity on average, with a mean error of 73 trips per day, although the median error is only -5 trips per day. This group exhibits the highest percentage errors, with a mean APE of 95 percent and a median APE of 49 percent. These elevated percentage errors are influenced by the relatively small number of observed counts at these locations, where even modest differences in estimated trips can produce large percentage deviations.

For sites with 150 to 500 average daily bicycle trips, StreetLight estimates align most closely with the observed counts. The mean error is -2 trips per day and the median error is -42 trips per day, indicating little systematic bias overall. The median APE of 37 percent is slightly lower than that observed for the lowest-volume sites, suggesting improved agreement when bicycle activity is moderate.

At sites with more than 500 average daily bicycle trips, StreetLight tends to report somewhat lower volumes than the counters, with a mean error of -119 trips per day and a median error of -213 trips per day. However, the percentage errors for this group are the lowest of the three volume categories. The mean APE is 34 percent and the median APE is 38 percent, indicating that although the absolute differences can be larger in terms of trips per day, the estimates remain relatively close to observed counts when viewed as a proportion of the total bicycle activity. Collectively, these results suggest that percentage agreement generally improves as observed bicycle volumes increase.

The facility-type comparison in Table 6 and Figure 23 indicates broadly similar performance across roadway and trail facilities. Road sites exhibit a mean error of -123 trips per day and a median error of -46 trips per day, while trail sites have a mean error of -8 trips per day and a median error of -16 trips per day. Although the trail sites display greater variability, including several large positive and negative deviations, the central tendency of the error distributions is similar between the two facility types. Likewise, the median APE is nearly identical for road and trail locations, at 39 percent and 38 percent, respectively. The higher mean APE observed for trails (62 percent versus 40 percent for roads) is influenced by a small number of large percentage-error outliers visible in Figure 23.

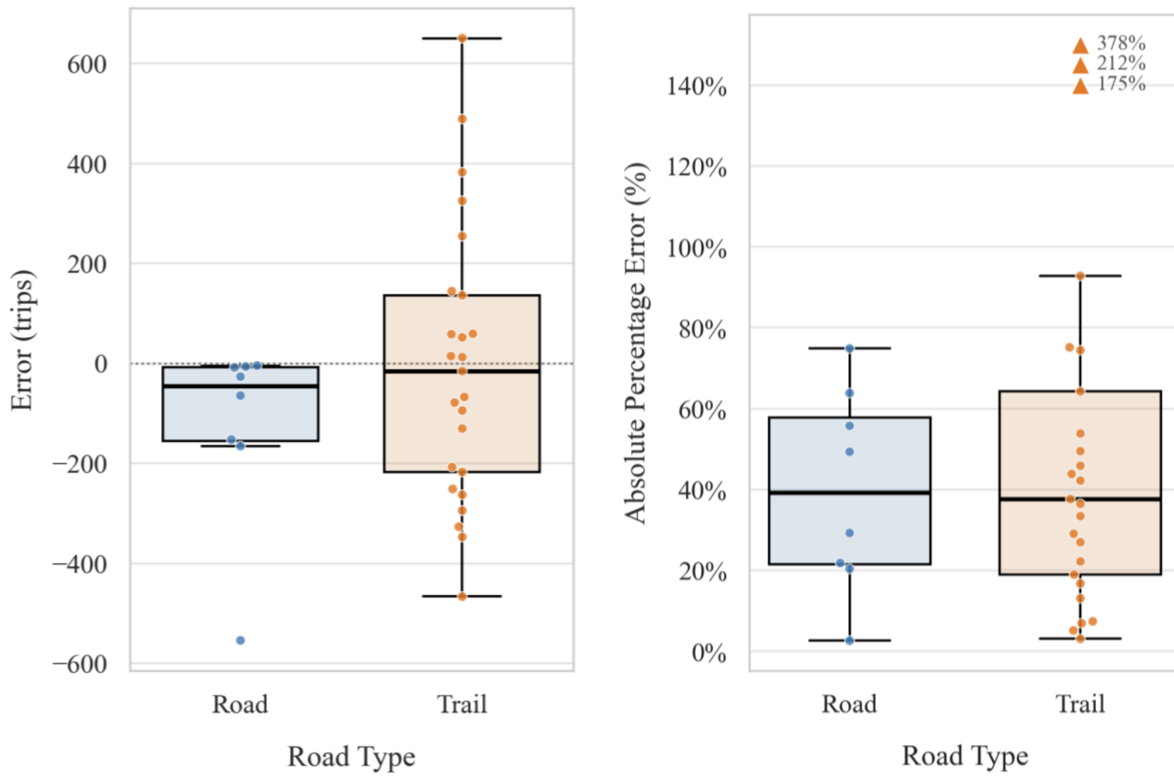

Figure 23 – Distribution of StreetLight errors and absolute percentage errors by facility type

It should be noted, however, that trail count locations generally exhibit higher observed bicycle volumes than roadway locations. The mean observed volume at trail sites is 523 trips per day, compared with 343 trips per day at roadway sites, while the corresponding median volumes are 467 and 177 trips per day. As a result, some differences in percentage-based error metrics may reflect underlying differences in activity levels rather than facility type alone. With a larger sample of count locations, it may be useful to evaluate facility type and volume effects simultaneously by comparing road and trail sites within common volume categories. However, such an analysis was not practical in the current study because subdividing the count locations by both facility type and volume category would result in very small sample sizes within each group, limiting the reliability of the resulting comparisons.

These results suggest that StreetLight provides reasonably consistent performance across different facility types and generally performs better at locations with higher bicycle activity. While the absolute difference between estimated and observed volumes can be larger at high-volume sites, the percentage errors are substantially lower, indicating stronger relative agreement with observed counts. Facility type also appears to have a limited effect on performance, with road and trail sites exhibiting similar median error and median APE values. These findings suggest that StreetLight may be useful for identifying relative bicycle activity patterns and broad demand trends, particularly on moderate- and high-volume facilities. However, local count data remains important when precise site-level volume estimates are required.

Illustrative StreetLight Count-Site Case Reviews

To complement the aggregate StreetLight count-site evaluation, a few cases were reviewed in greater detail. The examples are not intended to represent the full range of StreetLight performance, but they provide useful context for interpreting the aggregate error patterns in real facility settings. For each case, two figures are presented. The first figure documents the count-site context, including the counter location, shown by the red dot, and the nearby StreetLight segment representation, shown by the cyan line. This figure includes map, satellite, and Google Street View perspectives to show the counter location, the surrounding facility geometry, and the visible bicycle facility at the site. The second figure compares the StreetLight estimated volume with the corresponding counter-based volume and reports the resulting error and percentage error.

Case 1: Mount Vernon Trail Near 14th Street Bridge

Figure 24 shows the location of Flow Detectors 2380 and 2381 on the Mount Vernon Trail near the turnoff to 14th Street Bridge. The map and satellite views show that the counter is located on a major trail corridor near a relatively complex junction, where the Mount Vernon Trail continues south and another trail connection branches toward Boundary Drive. This configuration creates several nearby points where bicycle movements may split or merge. The Google Street View image confirms the off-street trail setting.

Figure 25 compares the StreetLight estimate with the observed counter volume for this case. StreetLight reports 1,179 annual average daily bicycle trips, compared with 1,043 trips from the counter. This corresponds to an error of 136 trips per day, or an overestimation of 13 percent. Despite this difference, the two values are relatively close given the magnitude of bicycle activity at the site. The figure also shows that StreetLight identifies this trail corridor as a high-volume bicycle facility, with segment-level volumes generally increasing toward the count location.

One segment immediately adjacent to the count site reports a substantially lower volume than neighboring segments, creating a localized discontinuity. However, this occurs near the complex trail junction described above, where bicycle trips may be redistributed across multiple nearby connections. Overall, the example demonstrates that StreetLight captures the corridor as a major bicycle facility and produces an estimate that is broadly consistent with the observed count, while also illustrating how complex trail geometry can affect interpretation of individual segment-level estimates.

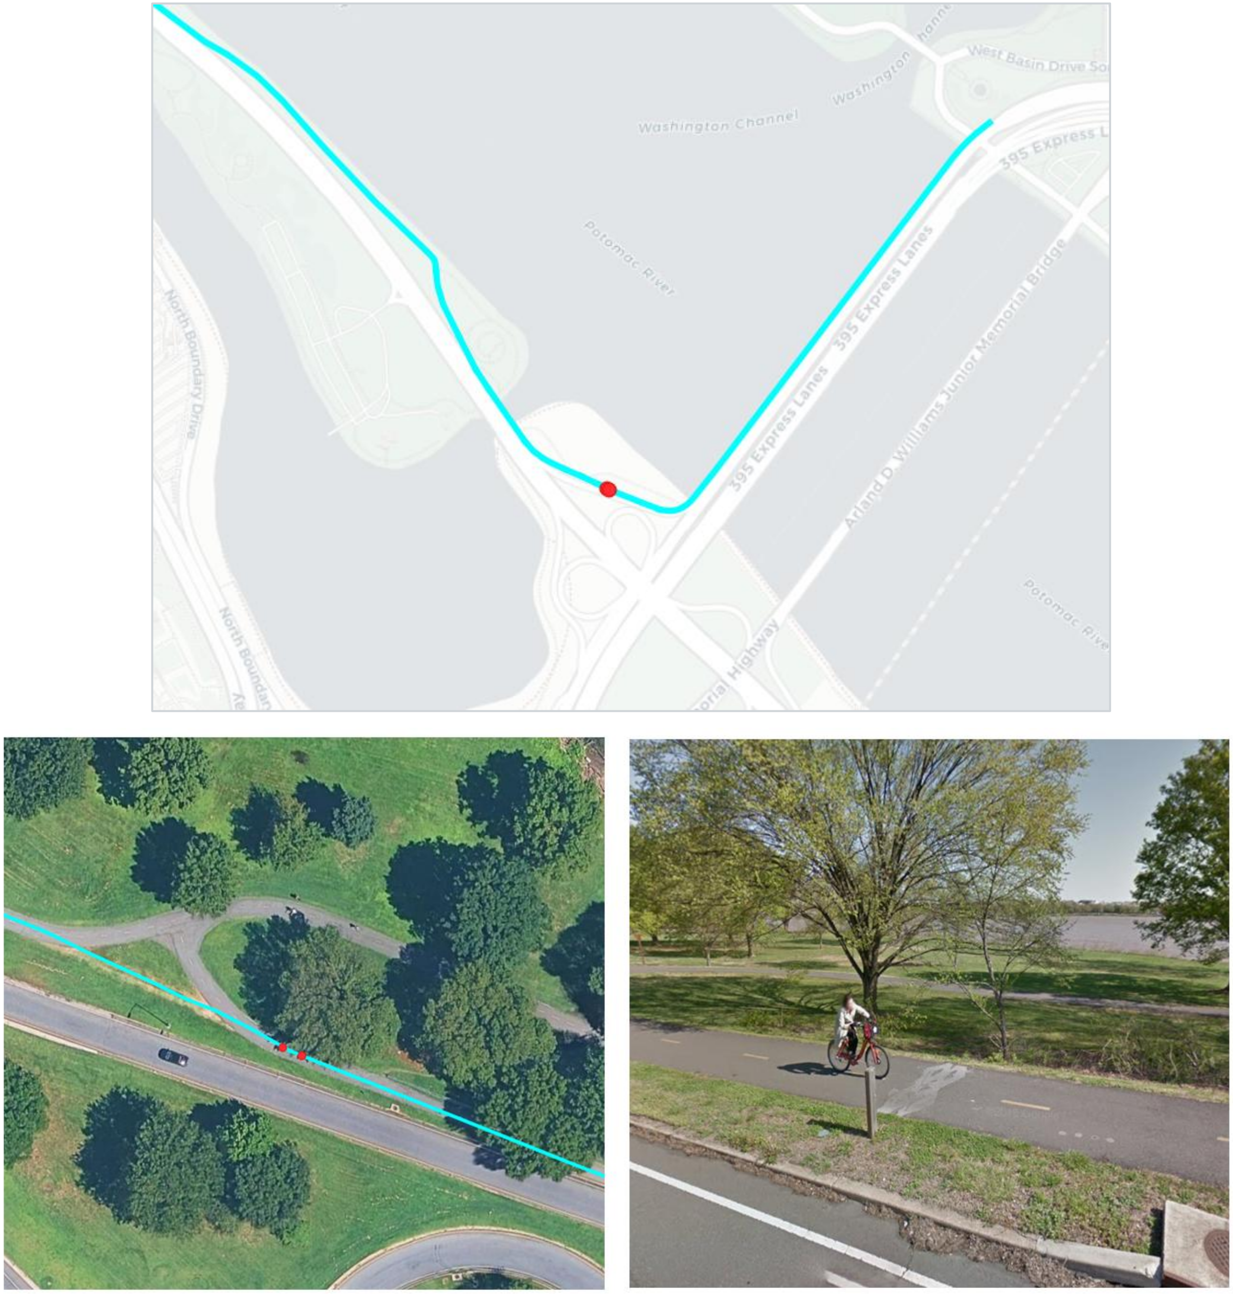

Figure 24 – StreetLight case 1 site context

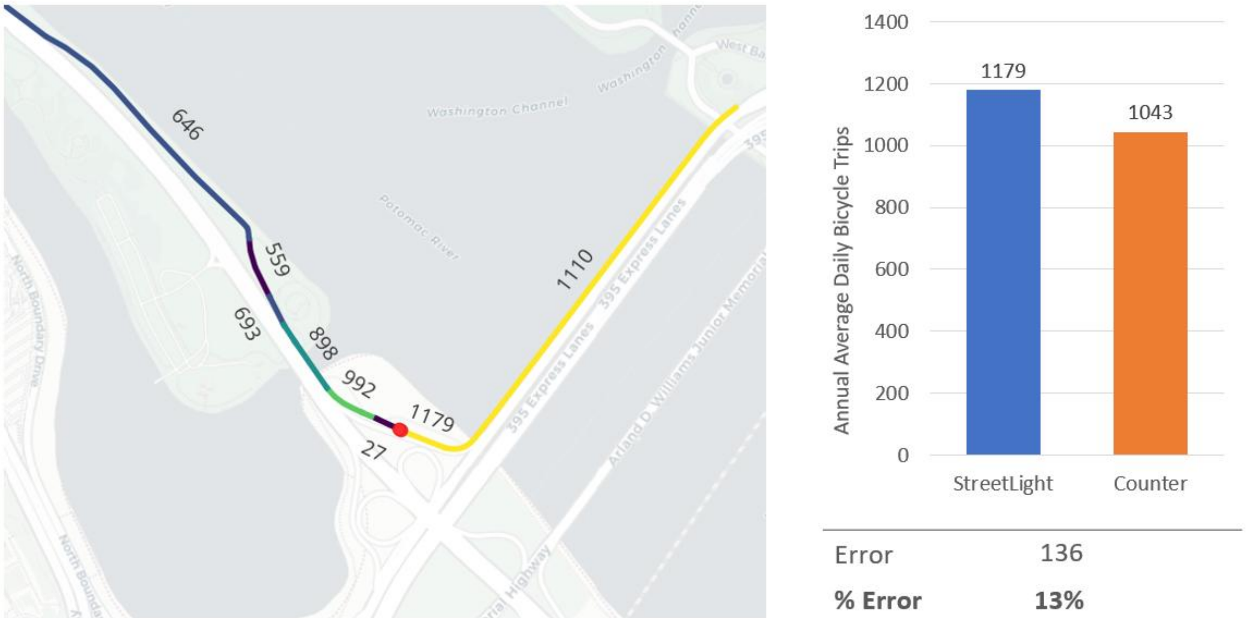

Figure 25 – StreetLight case 1 volume comparison

Case 2: Custis Trail in Rosslyn

Figure 26 shows the location of Flow Detectors 505 and 506 on the shared-use section of the Custis Trail in the Rosslyn urban village area, near Langston Boulevard and North Lynn Street in Arlington. The map, satellite image, and Google Street View image indicate that this is a trail facility in a dense and complex urban setting, with nearby roadway connections, intersections, and access points.

Figure 27 presents the volume comparison for this case. StreetLight reports 261 annual average daily bicycle trips, compared with 728 trips from the counter, resulting in an error of -467 trips per day and an underestimation of 64.2 percent. This represents a substantial difference between the vendor estimate and the observed count. The corridor-level pattern provides additional context. StreetLight identifies the facility as a relatively high-volume bicycle corridor, with adjacent segments generally reporting between approximately 190 and 380 bicycle trips per day. However, the count-site segment and nearby segments all remain well below the observed count, indicating that the underestimation is not limited to a single segment but extends across the surrounding corridor.

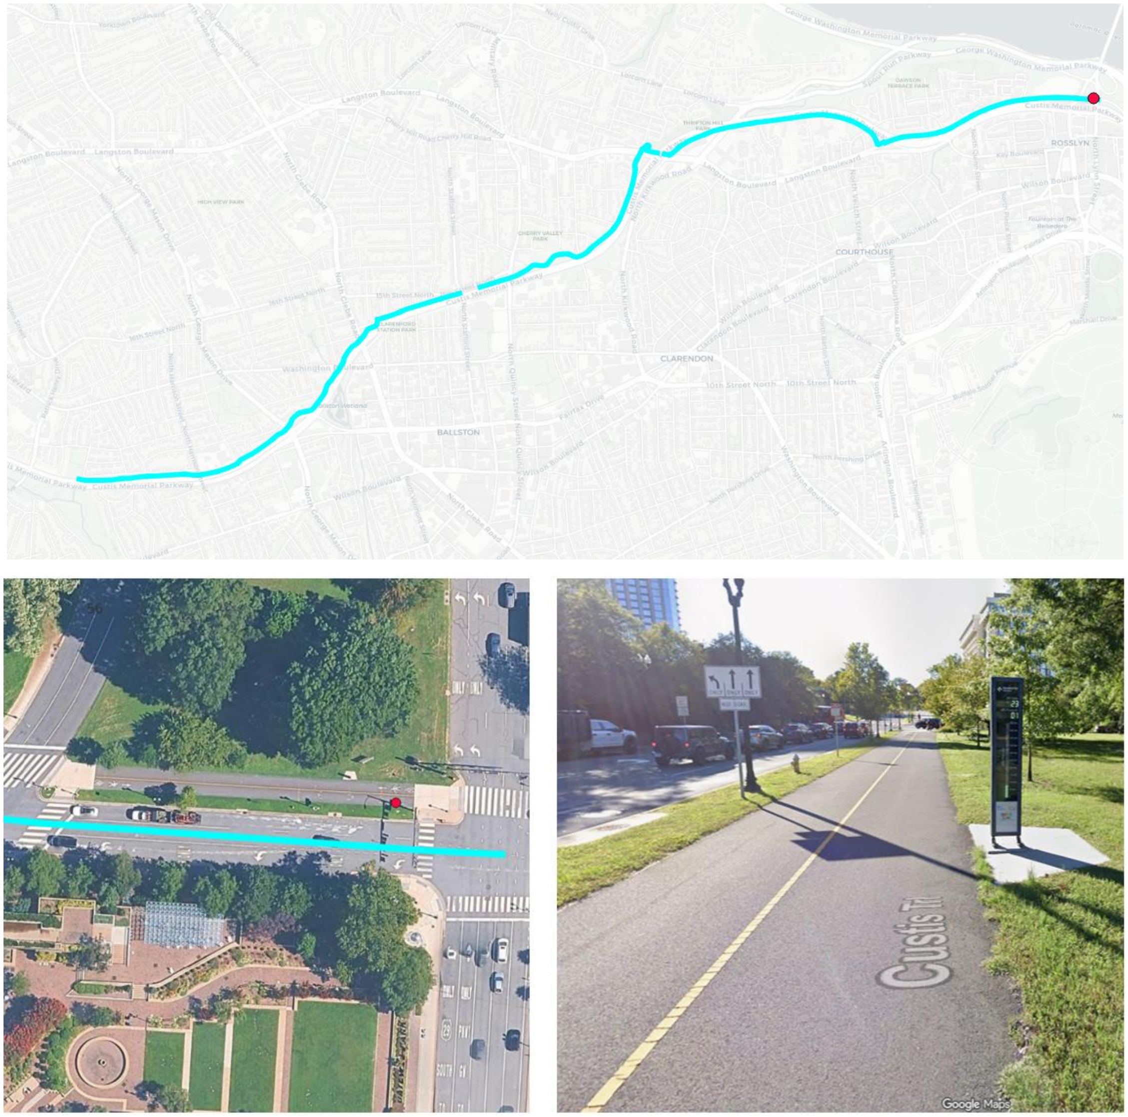

Figure 26 – StreetLight case 2 site context

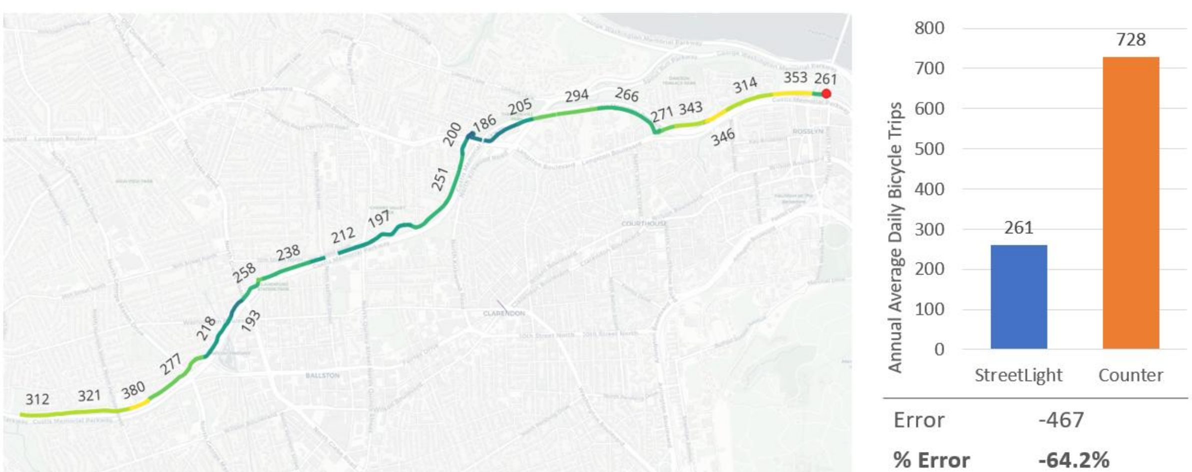

Figure 27 – StreetLight case 2 volume comparison

Case 3: Wilson Boulevard Bicycle Lane

Figure 28 shows the location of Flow Detector 335 on the westbound bicycle lane on the north side of Wilson Boulevard between North Danville Street and North Cleveland Street in Arlington. Unlike the previous two examples, this case represents an on-street urban bicycle facility rather than an off-street trail. However, interpretation of the observed volume at this site requires some caution. Cyclists traveling outside the bicycle lane, including those using the adjacent general-purpose lane, may not be detected. In addition, the counter does not identify direction of travel, so the analysis assumes that recorded bicycle activity corresponds to the westbound direction represented by the vendor segment.

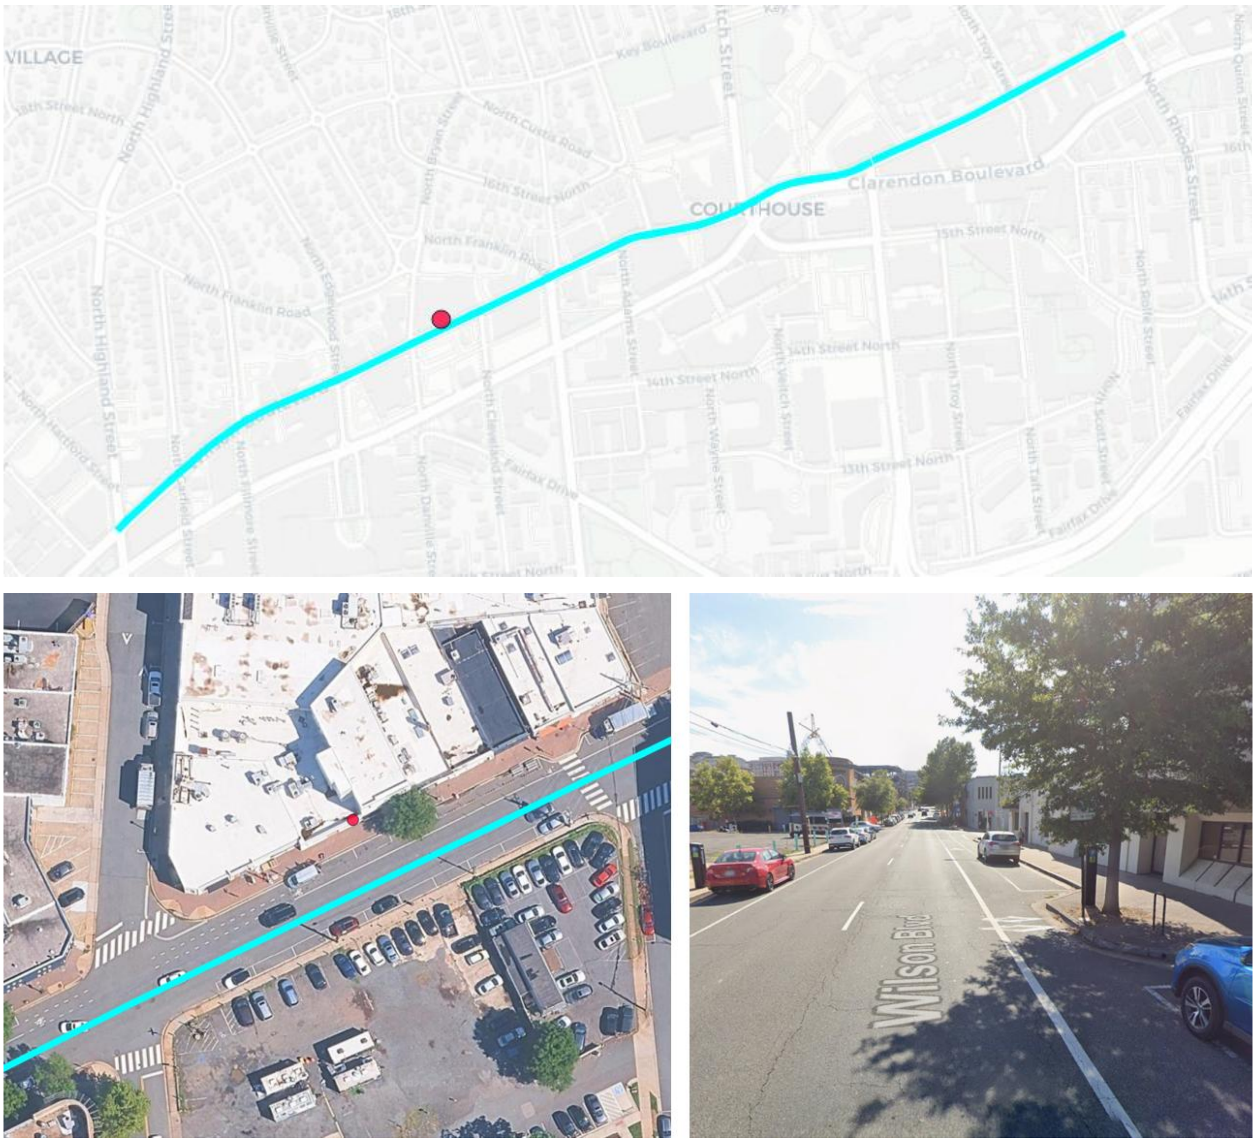

Figure 28 – StreetLight case 3 site context

Figure 29 compares the StreetLight estimate with the counter-based volume for this location. StreetLight reports 56 annual average daily bicycle trips, compared with 222 trips from the counter, resulting in an error of -166 trips per day and an underestimation of 74.8 percent. This case demonstrates that substantial underestimation is not limited to trail facilities. Even on an urban roadway bicycle facility, StreetLight can report volumes that are considerably lower than observed counts. The surrounding segments also exhibit relatively similar estimated volumes, suggesting that the underestimation reflects a broader corridor-level pattern rather than an isolated segment anomaly.

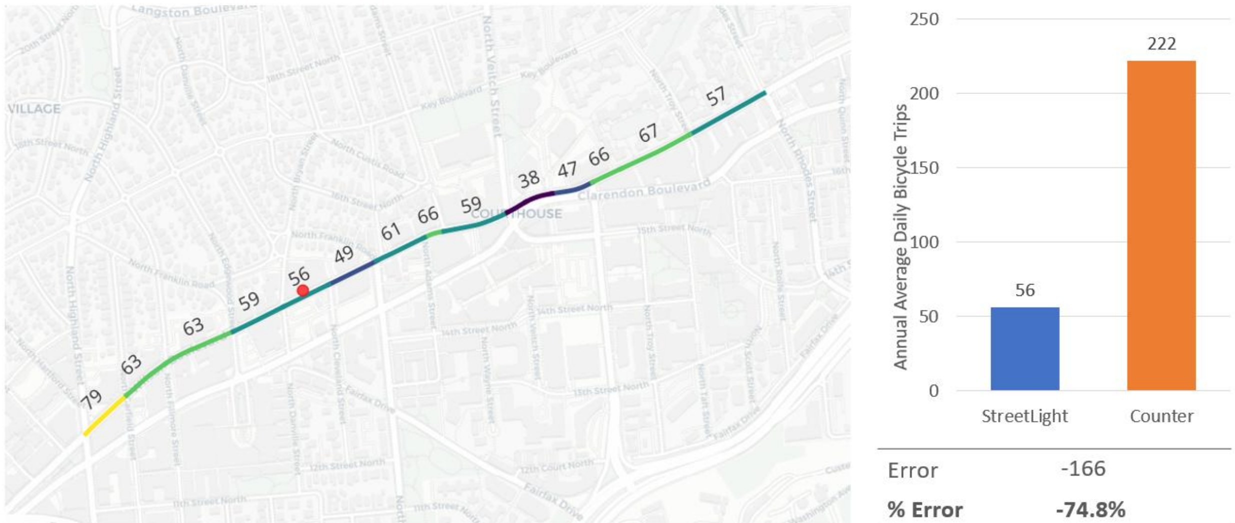

Figure 29 – StreetLight case 3 volume comparison

Together, these case reviews provide a site-level interpretation of the broader StreetLight findings. While the aggregate results indicate generally reasonable agreement with observed counts across the full set of locations, individual sites can still exhibit substantial deviations. The three examples illustrate the range of outcomes observed in the validation dataset: one high-volume trail facility with relatively close agreement, one trail facility with substantial underestimation, and one roadway bicycle facility with substantial underestimation. These examples highlight that StreetLight can identify important bicycle corridors and capture relative activity patterns, while site-level volume estimates may vary considerably from observed counts in some locations.

Replica Count-Site Comparison

For Replica, the count-site comparison includes 68 directional comparison cases, reflecting Replica's directional segment reporting and the corresponding directional counter volumes. Figure 30 compares counter-based average daily bicycle trips with Replica's reported directional average daily bicycle trips. The Spearman correlation of 0.55 indicates a moderate positive relationship; locations with higher observed bicycle volumes generally tend to have higher Replica estimates, but the site-level agreement is variable. The fitted regression line has a slope of 0.68, suggesting that Replica captures some of the increasing trend in bicycle volume but does not follow the 1:1 relationship closely across the full range of count sites. Compared with the counters, Replica frequently reports higher values, especially for many low-volume locations. This is reflected in the mean absolute percentage error of 131 percent, the mean absolute error of 280 trips per day, and the root mean square error of 399 trips per day, which indicate substantial differences at individual count locations.

Figure 31 provides additional insight into the direction and structure of these differences by plotting Replica minus counter volume against the observed counter volume. Replica shows a positive mean bias of 206 trips per day, indicating that Replica reports higher volumes than the counters on average across the comparison sites. The LOESS trend line is generally above zero for much of the observed volume range, although the differences are not uniform across all sites. Some locations show underestimation, particularly in the middle of the count-volume range, but many sites show positive errors, and several high-volume locations have large overestimates. The proportional bias statistic is positive and statistically significant, suggesting that Replica's error tends to increase s as observed counter volume increases.

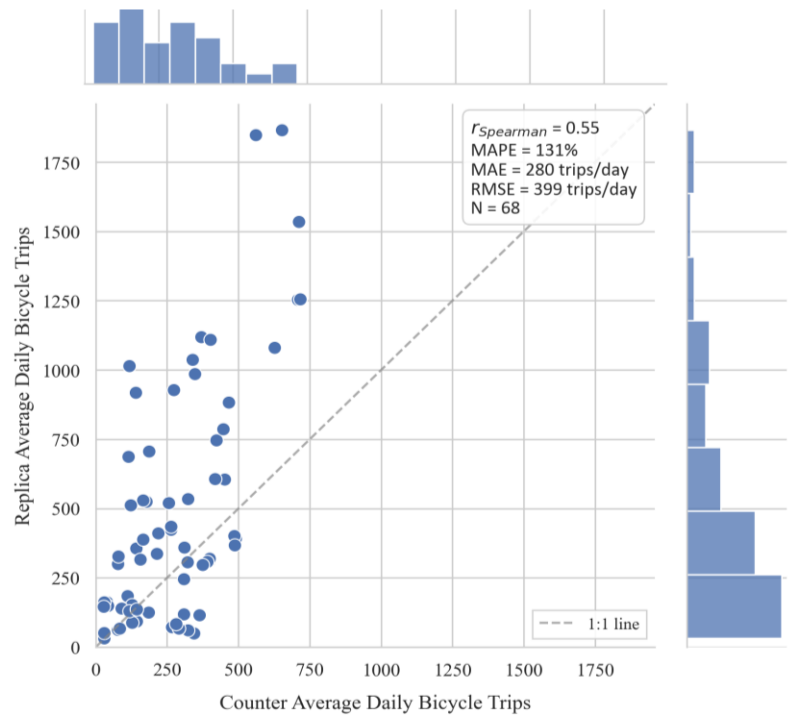

Figure 30 – Comparison of Replica and agency count average daily bicycle volumes

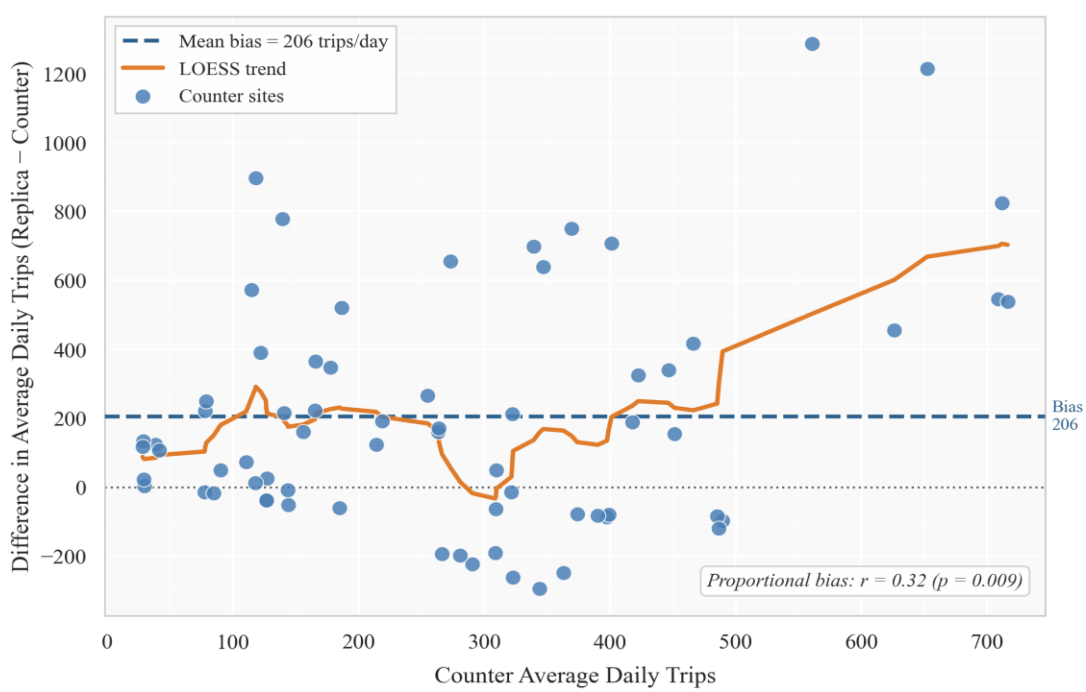

Figure 31 – Difference between Replica and agency count volumes by count-site volume

Table 7, Figure 32, and Figure 33 examine how Replica errors vary by observed bicycle volume range and facility type. Error is defined as Replica minus counter volume, so positive values indicate that Replica reports higher volumes than the agency counter, while negative values indicate underestimation. Absolute percentage error (APE) measures the size of the difference relative to the counter volume, regardless of direction.

Because Replica reports directional segment volumes, the count-volume ranges used for this analysis differ from the StreetLight bins. For StreetLight, the volume ranges were based on bidirectional counter volumes, using thresholds of 150 and 500 average daily bicycle trips. For Replica, these thresholds were adjusted to reflect the directional structure of the comparison. A direct directional adjustment would suggest approximately halving the StreetLight thresholds to 75 and 250 trips per day. However, Replica is evaluated using fall 2024 ADT values, while StreetLight is evaluated using annual average daily volumes. Since fall bicycle volumes are typically higher than annual averages, the thresholds were rounded upward to 100 and 300 trips per day. As a result, the Replica volume bins are defined as less than 100, 100 to 300, and greater than 300 directional average daily bicycle trips.

The corresponding boxplots summarize the distribution of errors and APE values within each group. As with the StreetLight analysis, the box represents the middle 50 percent of observations, from the 25th to the 75th percentile, with the median shown as the line inside the box. The whiskers show the estimated non-outlier range, and the overlaid dots show the individual comparison cases directly. The P25, P50, and P75 values in Table 7 correspond to the same percentile values shown in the boxplots, allowing the table and figures to be interpreted together.

Table 7 – Replica error summary by counter volume range and facility type

Scenario		N	Error (trips)				Absolute Percentage Error (%)			
			Mean	P25	P50	P75	Mean	P25	P50	P75
Overall		68	206	-41	128	371	131%	30%	76%	188%
By Volume Range	<100	11	91	13	108	128	204%	37%	260%	317%
	100-300	27	198	-23	161	356	159%	46%	77%	208%
	>300	30	255	-84	172	543	80%	21%	70%	88%
By Facility Type	Road	19	141	6	118	214	142%	40%	70%	179%
	Trail	49	231	-51	155	417	127%	30%	76%	186%

N: Number of cases; P25: 25th percentile; P50: 50th percentile (Median); P75: 75th percentile

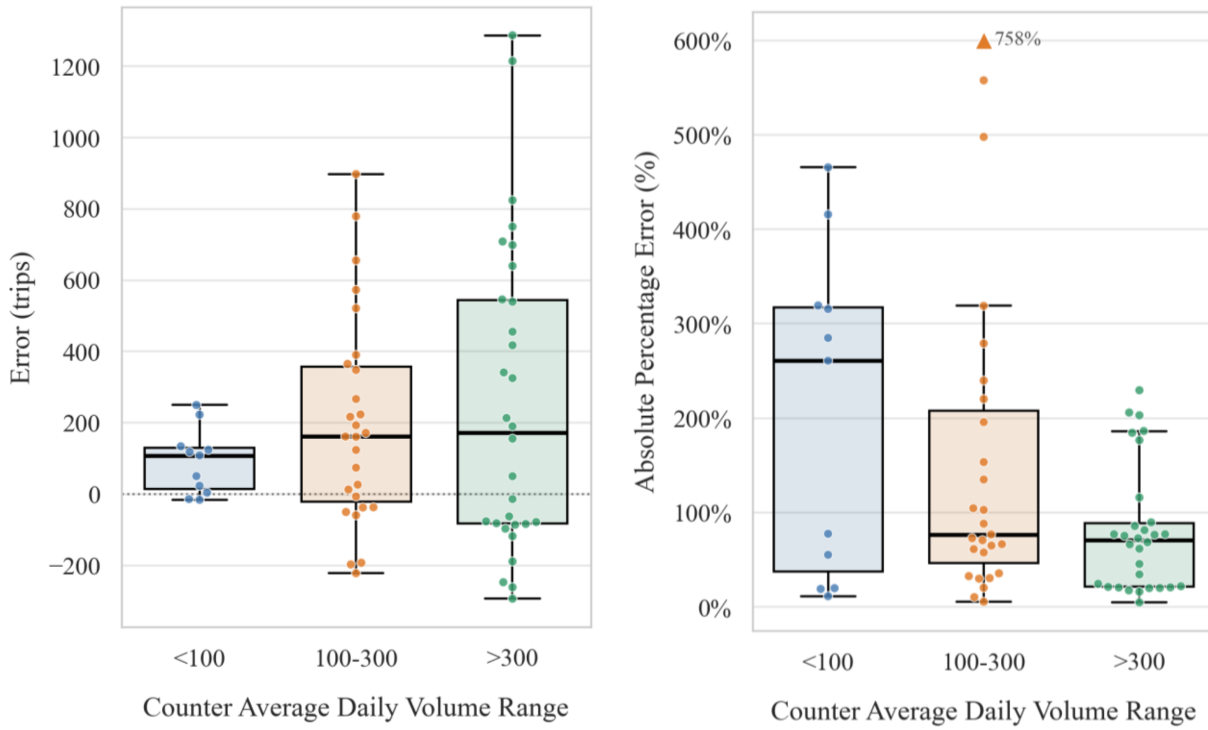

Figure 32 – Distribution of Replica errors and absolute percentage errors by counter volume range

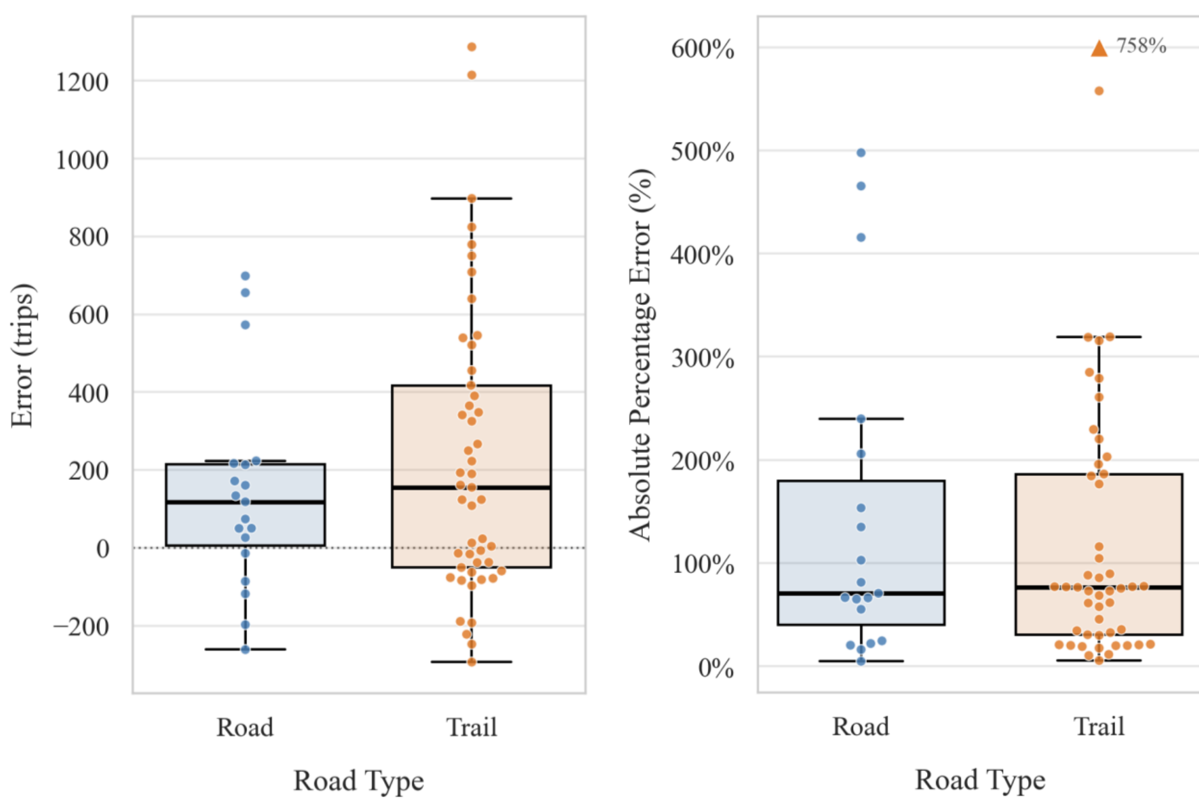

Figure 33 – Distribution of Replica errors and absolute percentage errors by facility type

The lowest-volume sites show relatively compact trip-level errors but very large percentage errors. For sites below 100 directional average daily bicycle trips, the middle 50 percent of errors ranges from 13 to 128 trips per day, with a median error of 108 trips per day. This indicates that most low-volume sites are overestimated in absolute terms, although the size of the trip-level errors is smaller than in the higher-volume bins. However, because the observed counts are small, these differences translate into very large APE values. The median APE is 260 percent, meaning that the typical relative error in this bin is large even when the trip difference is not in absolute terms.

For the middle-volume sites, between 100 and 300 directional average daily bicycle trips, the boxplots show a wider and more mixed error distribution. The median error is 161 trips per day, but the 25th percentile is slightly negative, indicating that some sites are underestimated while many others are overestimated. The APE distribution is also highly skewed, the median APE is 77 percent, while the mean APE is 159 percent. This difference between the mean and median reflects the influence of large percentage-error outliers, including the extreme APE value visible in the boxplot. For this group, the median provides a more stable indication of typical relative error, while the mean reflects the effect of a small number of very large discrepancies.

For the highest-volume sites, above 300 directional average daily bicycle trips, the absolute errors are larger and more dispersed, which is expected to some extent because the observed volumes are larger. The median error is 172 trips per day, and the 75th percentile reaches 543 trips per day, indicating that many high-volume sites are overestimated, with some cases differing from the counters by several hundred trips per day. At the same time, the APE values are lower than in the other two volume bins, with a median of 70 percent and a mean of 80 percent. This does not necessarily indicate stronger agreement in practical terms; rather, it reflects that larger counter volumes reduce the percentage impact of a given trip-level difference. The high-volume bin therefore shows the importance of interpreting trip-level errors and APE together.

The facility-type comparison shows a less distinct separation than the volume-bin comparison. Both road and trail sites have positive median errors, indicating that Replica tends to overestimate volumes for both facility types. However, the magnitude and spread of trip-level errors are somewhat larger for trails. Trail sites have a median error of 155 trips per day, compared with 118 trips per day for road sites, and the upper quartile reaches 417 trips per day for trails compared with 214 trips per day for roads. The trail boxplot also shows a wider distribution, suggesting greater site-to-site variability. At the same time, the 25th percentile error for trails is negative, indicating that some trail locations are underestimated despite the overall tendency toward positive errors. In percentage terms, the distinction between road and trail sites is less pronounced: median APE values are similar, at 70 percent for roads and 76 percent for trails, and the interquartile ranges overlap substantially. Facility type appears to contribute to differences in trip-level error magnitude, but it does not explain Replica’s error pattern as clearly as observed volume.

It should be noted that the observed bicycle volumes at road and trail count locations are relatively similar in the Replica sample. Roadway sites have a mean observed volume of 225 trips per day and a median of 264 trips per day, while trail sites have a mean observed volume of 296 trips per day and a median of 266 trips per day. As a result, differences in the error distributions between facility types are less likely to reflect underlying differences in observed volume than in the StreetLight analysis.

As a sensitivity check, the Replica directional volumes were also aggregated to bidirectional volumes and compared with corresponding bidirectional count-site observations. The bidirectional comparison produced similar error patterns and did not materially change the overall findings. However, the Spearman correlation increased from 0.55 to 0.70, suggesting that directional over- and under-estimation may partially offset one another in some cases. Because the broader interpretation of the results remained unchanged, only the directional analysis is presented in detail.

These results indicate that Replica estimates should also be interpreted carefully for site-level bicycle volume analysis, but for somewhat different reasons than StreetLight. Replica tends to report higher volumes than the counters in many cases, with substantial site-to-site variability. The results suggest that low-volume locations may be overstated in relative terms, while higher-volume locations can show large absolute differences in daily bicycle trips even when percentage errors appear lower. As with StreetLight, Replica may provide useful information about relative activity patterns and corridor-level demand, but local count data or calibration would be important before using the estimates as direct substitutes for observed volumes in applications that depend on accurate site-level magnitude.

Illustrative Replica Count-Site Case Reviews

Following the aggregate Replica count-site evaluation, several individual locations were reviewed more closely to provide site-level context for the statistical results. These examples illustrate how Replica's directional segment estimates compare with corresponding directional count volumes under different facility and network conditions. They should be interpreted as illustrative cases rather than as a representative sample of all Replica comparison locations.

Each case is presented using two figures. The first figure provides the spatial and physical context of the count site, with the counter location identified by a red dot and the nearby Replica segment shown in cyan. The map, satellite image, and Google Street View panel show how the counter relates to the surrounding roadway or trail facility and to Replica's segment representation. The second figure presents the volume comparison for the same location, including the Replica directional estimate, the corresponding counter-based directional volume, and the resulting error and percentage error. This format allows the numerical difference to be interpreted alongside the local facility context and the specific segment used for the comparison.

Case 1: Washington & Old Dominion Trail Near Columbia Pike

Figure 34 shows the location of Flow Detectors 493 and 495 on the shared-use section of the Washington & Old Dominion Trail near Columbia Pike in Arlington. The map, satellite image, and Google Street View panel provide context for the trail setting and show how the counter location relates to Replica's directional segment representation.

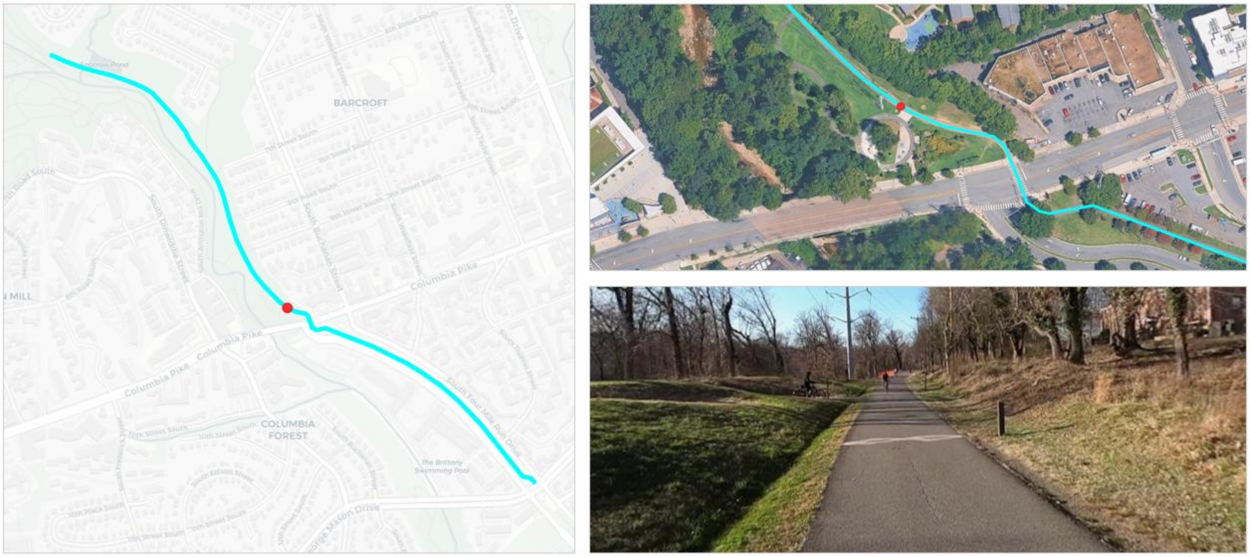

Figure 34 – Replica case 1 site context

Figure 35 presents the directional volume comparison for this case. Replica reports 320 average daily bicycle trips in the northbound direction compared with 399 trips from the counter, corresponding to an error of -79 trips per day, or -20 percent. In the southbound direction, Replica reports 308 trips compared with 390 trips from the counter, corresponding to an error of -82 trips per day, or -21 percent. This is one of the clearer examples of directional consistency between Replica and the observed counts and the magnitude of underestimation is relatively limited. The case indicates that, under some trail conditions, Replica can preserve both the directional balance and approximate scale of observed bicycle activity.

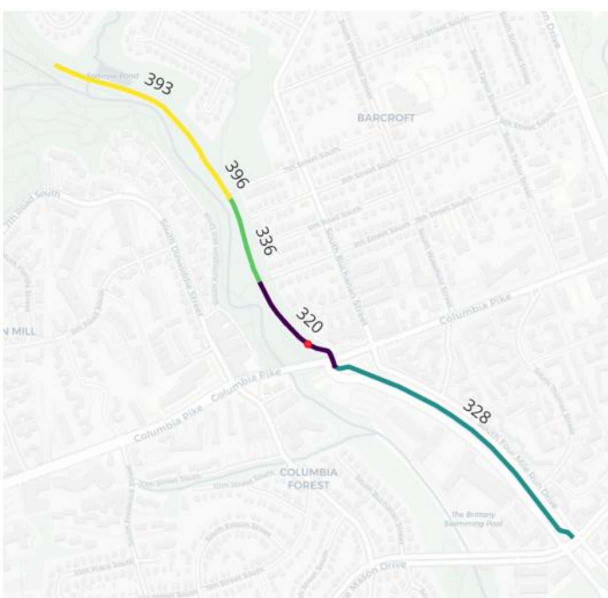

(a) Northbound segment volumes

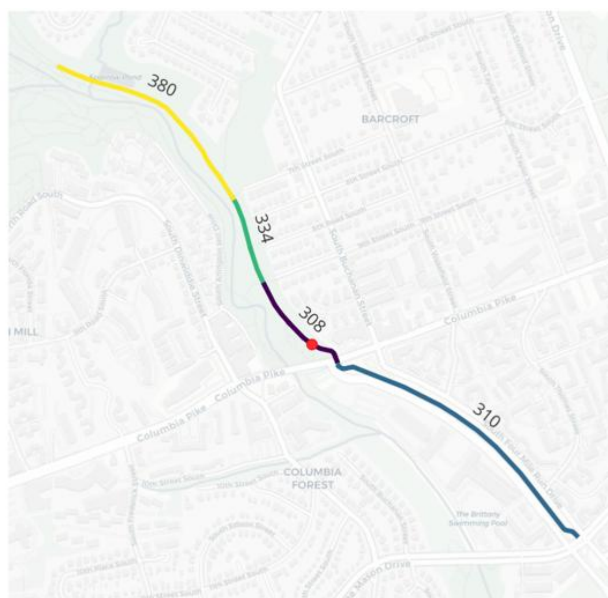

(b) Southbound segment volumes

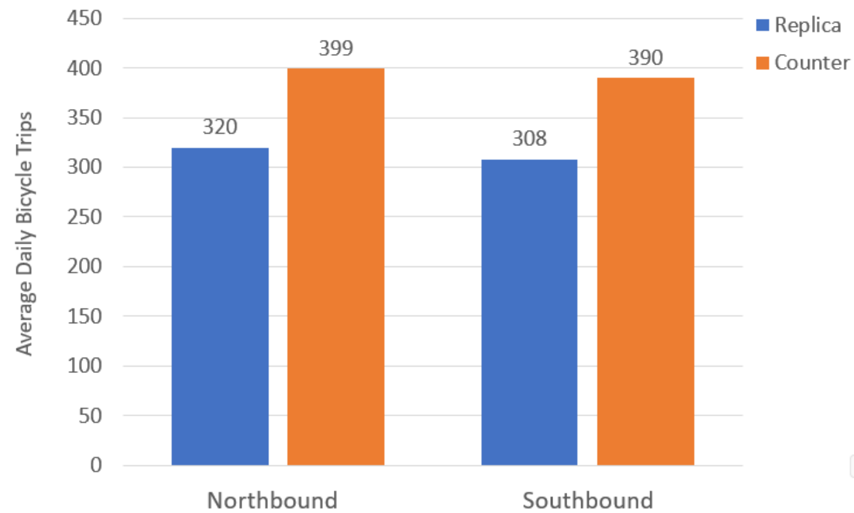

	Northbound	Southbound
Error	-79	-82
Error %	-20%	-21%

(c) Replica vs. counter volumes

Figure 35 – Replica case 1 directional volume comparison

Case 2: Mount Vernon Trail Near 14th Street Bridge

Figure 36 shows the location of Flow Detectors 2380 and 2381 on the Mount Vernon Trail near the turnoff to 14th Street Bridge. The map, satellite image, and Google Street View panel indicate that the counter is located on a major trail corridor near a relatively complex junction, where the Mount Vernon Trail continues south and another trail connection branches toward Boundary Drive.

Figure 36 – Replica case 2 site context

Figure 37 compares Replica's directional estimates with the observed counter volumes. Replica reports 1,847 average daily bicycle trips in one direction compared with 561 trips from the counter, producing an error of 1,286 trips per day, or 229 percent. In the opposite direction, Replica reports 1,866 trips compared with 652 trips from the counter, producing an error of 1,214 trips per day, or 186 percent. The directional balance in the Replica estimates is similar to the counter pattern, but the overall magnitude is much higher than the observed volumes. The case illustrates that a vendor estimate may reproduce directional symmetry while still substantially overstating total activity at a specific trail location. However, Replica volumes on adjacent segments west of the count site are more consistent with the observed counts, indicating that agreement may vary considerably over short distances along the corridor.

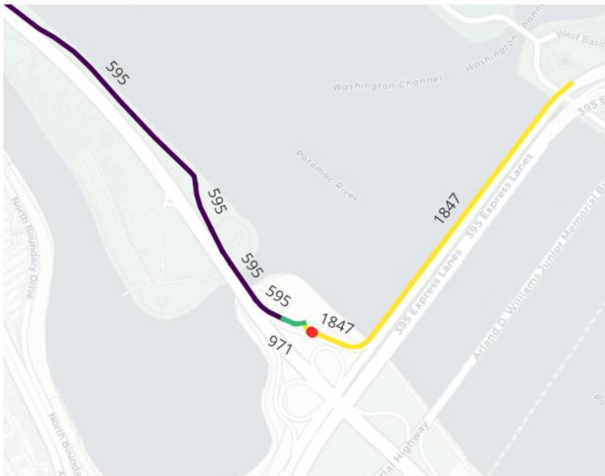

(a) Eastbound segment volumes

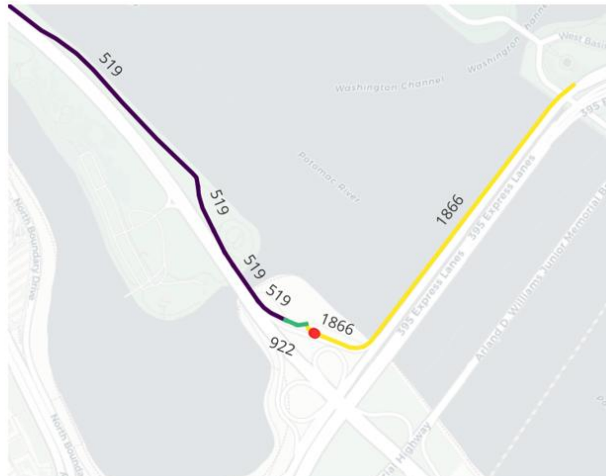

(b) Westbound segment volumes

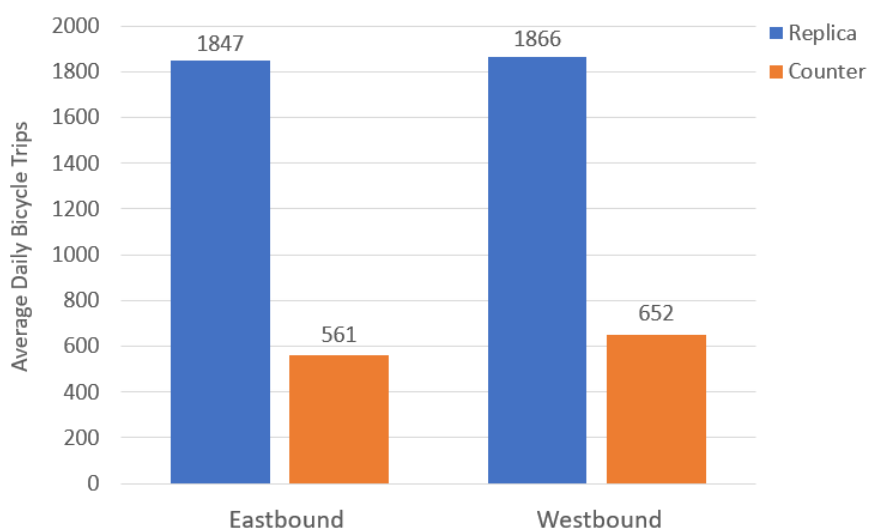

	Eastbound	Westbound
Error	1286	1214
Error %	229%	186%

(c) Replica vs. counter volumes

Figure 37 – Replica case 2 directional volume comparison

Case 3: Custis Trail in Rosslyn

Figure 38 shows the location of Flow Detectors 505 and 506 on the shared-use section of the Custis Trail in the Rosslyn urban village area, adjacent to Langston Boulevard and North Lynn Street in Arlington. The map, satellite image, and Google Street View panel document the counter location, the surrounding facility setting, and the trail segments included in the review.

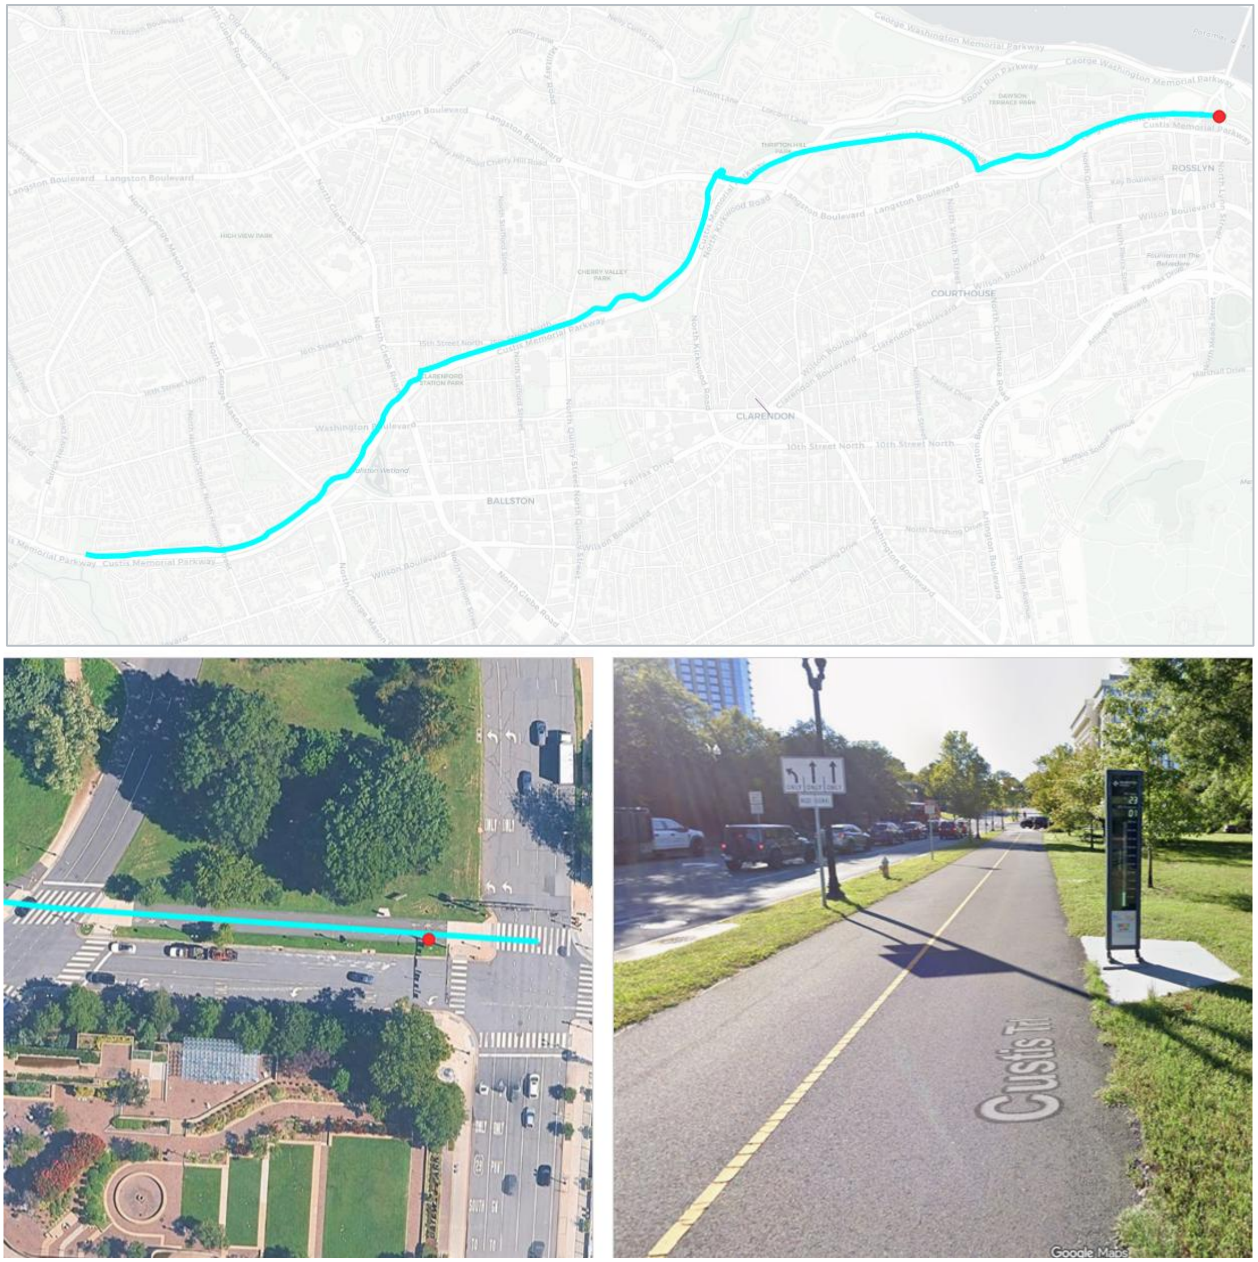

Figure 38 – Replica case 3 site context

Figure 39 presents the volume comparison for this case. Replica reports higher volumes than the observed counts in both directions, with an error of 417 trips per day in the eastbound direction and 749 trips per day in the westbound direction, corresponding to percentage errors of 89 percent and 203 percent, respectively. The directional pattern also differs from the counter observations: the counters show higher volume in the eastbound direction, while Replica reports higher volume in the opposite direction. The corridor-level view also shows that Replica estimates are not uniform along the selected Custis Trail segments. Adjacent segments have substantially higher estimated volumes in both directions, nearly twice the volume assigned to the matched segment. This suggests that the difference at the counter location may be related not only to the overall scale of Replica's estimates, but also to local variation in how bicycle activity is distributed across nearby directional trail segments.

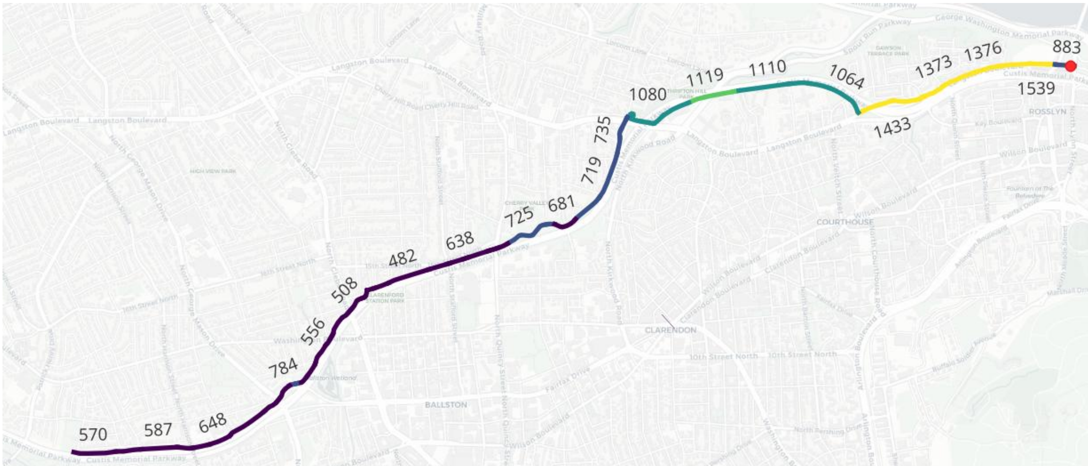

(a) Eastbound segment volumes

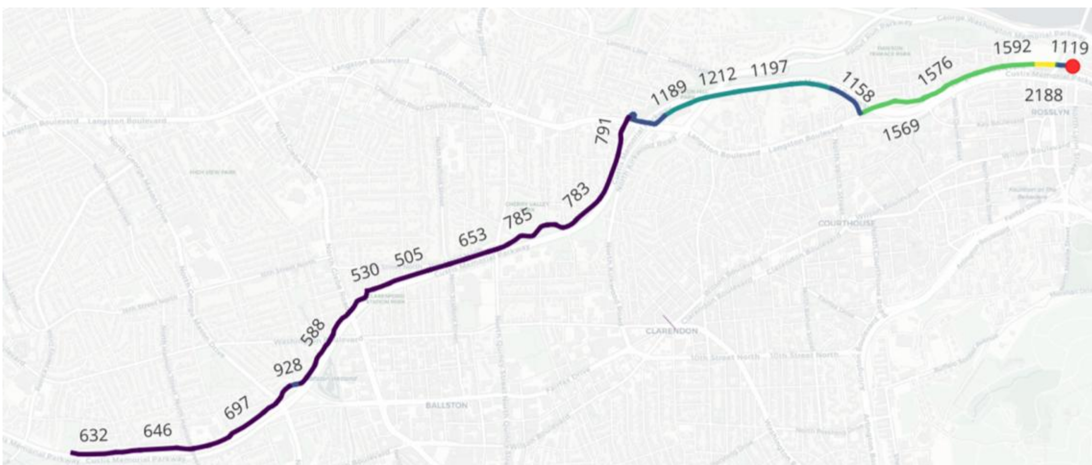

(b) Westbound segment volumes

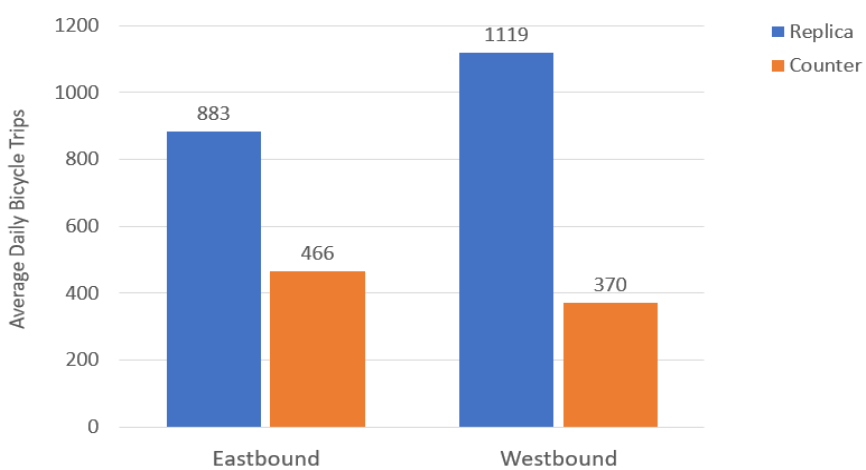

	Eastbound	Westbound
Error	417	749
Error %	89%	203%

(c) Replica vs. counter volumes

Figure 39 – Replica case 3 directional volume comparison

Case 4: East Capitol Street NE

Figure 40 shows the location of Flow Detector 2848 on East Capitol Street NE between 5th and 6th Streets SE in Washington, DC. Unlike the previous trail examples, this case represents an urban roadway bicycle facility. The map, satellite image, and Google Street View panel provide context for the street setting and the bicycle facility where the counter is located.

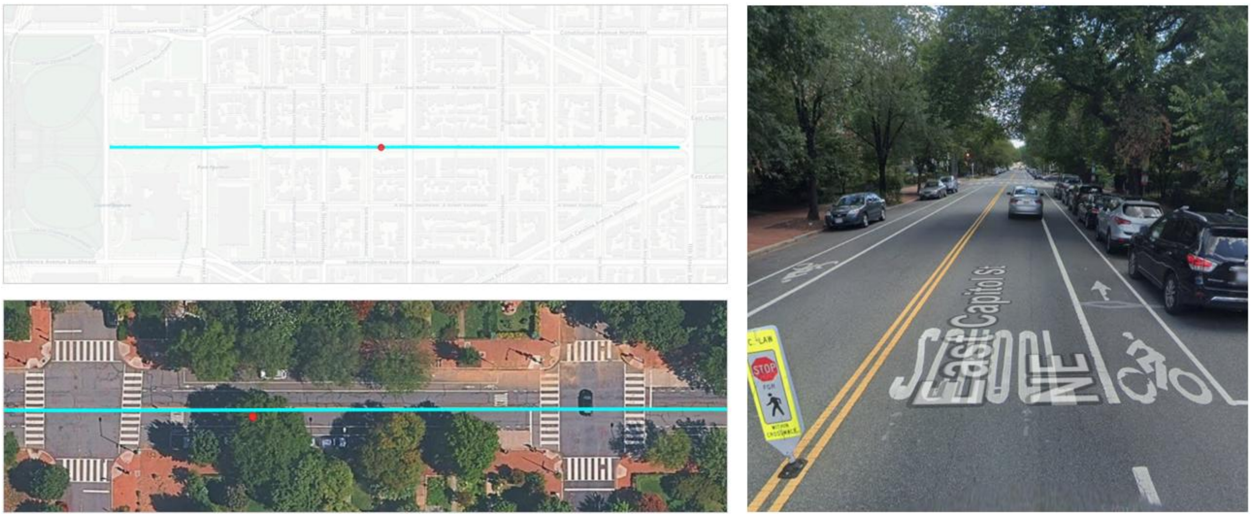

Figure 40 – Replica case 4 site context

Figure 41 shows Replica estimates for both directions on the nearby roadway segment; however, the count site includes only an eastbound counter and does not include a westbound counter. Therefore, the quantitative comparison for this case is limited to the eastbound direction. In that direction, Replica reports 61 average daily bicycle trips, compared with 323 trips from the counter. This corresponds to an error of -262 trips per day, or an underestimation of 81 percent. This case provides a useful counterpoint to the broader Replica results because it shows that the overall tendency toward overestimation does not apply uniformly across all sites. At this roadway location, Replica assigns substantially fewer bicycle trips than observed by the eastbound counter, highlighting the importance of interpreting directional segment estimates in relation to the specific counter configuration and facility context.

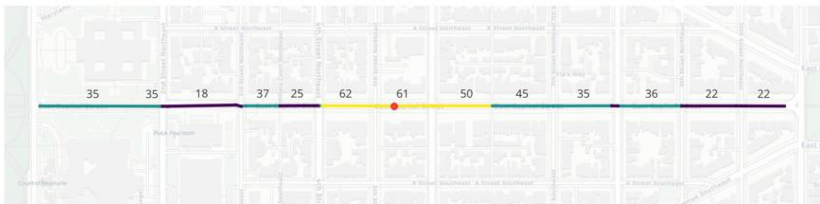

(a) Eastbound segment volumes

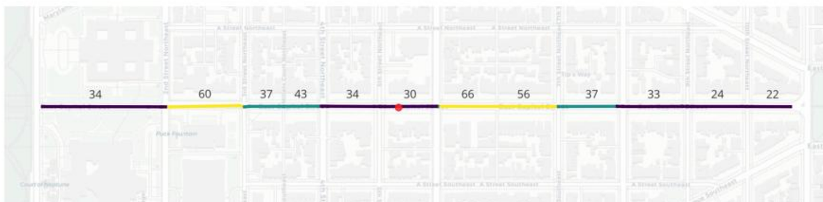

(b) Westbound segment volumes

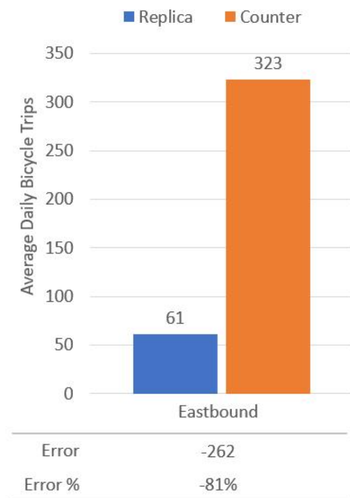

(c) Replica vs. counter volumes

Figure 41 – Replica case 4 directional volume comparison

Taken together, the Replica results show that the dataset provides useful directional and corridor-level information, but site-level volume interpretation requires caution. The aggregate analysis and case reviews both indicate that Replica estimates can differ substantially from observed counts, with errors varying by location, direction, and local segment context. The illustrative cases are especially useful because they show that differences are not always explained by a single factor, such as facility type or overall volume level. Instead, the comparison depends on how directional activity is assigned to individual segments and how well those segments correspond to the specific counter location.

Analysis #3 (Segment-Level Evaluation at Agency Count Sites) Summary

- **StreetLight:** Bi-Directional volume estimates showed strong rank-order agreement with agency counts and relatively little average bias, although moderate errors and substantial differences at some individual locations remained.
- **Replica:** Directional volume estimates showed moderate directional rank-order agreement and generally exceeded observed counts, with greater site-to-site variability. Agreement improved when directional volumes were aggregated, but the overall interpretation remained.

These segment-level products may be useful for characterizing relative bicycle activity and corridor-level demand but should not be treated as a direct substitute for local counts when accurate site-level volumes are required. Local validation or calibration remains important, particularly because performance can vary by location and segment configuration.

4. Conclusions

This bike/ped data validation report focuses on two components: (I) a product characterization matrix intended to help explain the nuances of each vendor's bike/ped datasets sold through the TDM, and (II) a case study in the DC metropolitan area. Given that this report marks the first time acquiring commercial bike/ped data, the case study focuses on several fundamental comparisons identified by the Technical Advisory Committee as being particularly important. The sample datasets were chosen to coincide with other complementary datasets, including agency-managed bike/ped counters in DC and Arlington, publicly available Capital Bikeshare trips data, the American Community Survey, and sample data provided by Strava Metro. Several key takeaways emerged:

- **Crowdsourced bike/ped products are diverse, making direct comparisons and evaluation challenging.** Although each of the five vendors offers data products that are categorized as "bike/ped", the datasets are distinct and quantify bike and pedestrian activity using different perspectives. For example, the bike/ped datasets vary in spatial resolution, with some vendors providing segment-level volumes (or normalized indices) and others providing zone-level trip estimates – either as trip volumes within an area or Origin-Destination flows or between geographies. Other differences include level of temporal aggregation, whether measures of activity represent samples of the population or population-level estimates. Additionally, particularly for zone-level results, reference/ground truth datasets are difficult – if not impossible – to acquire.

  Part I of this report provides a detailed characterization of all TDM bike/ped products to help provide clarity on each vendor's offerings.

- **Industry is in early product stage maturity but evolving quickly.** Five vendors in the TDM currently offer some type of scalable "bike/ped" product, and these offerings continue to evolve quickly. During the course of this validation activity, one vendor released an updated version of their model that substantially changed volume estimates relative to the initial version they had submitted. Another vendor announced the release of a new segment-level product that had not previously existed during the initial scan of TDM bike/ped products. This quickly changing landscape is encouraging, as it indicates that industry is working to bring better bike/ped data to market. However, when products are regularly updated, it can be difficult to use data for tracking trends over time, as changes in reported trips may be attributable to methodological differences rather than changes in user behavior.

- **Fundamental comparisons from the DC area case study yielded several insights at the zone and segment level:**
  - **Zone-Level (Bike):** Vendor datasets showed similar relative spatial patterns to one another and to a comparative dataset, although the magnitude of trip counts differed across vendors.
  - **Zone-Level (Pedestrian):** Vendor datasets generally identified central Washington, DC as a major concentration of pedestrian activity, but differed in total reported volume and in the spatial distribution of activity outside the urban core.
  - **Segment Level (Bike):** On a network scale, vendors showed varying spatial patterns, with only moderate agreement in which segments were associated with the highest bicycle activity. Follow-up analysis evaluated two vendors' bicycle volume estimates at agency count sites, providing an initial characterization of product accuracy on both roads and trails.

Note that this case study focused specifically on the DC metro area, which has high levels of bike/ped activity. Further work is needed to understand whether the findings from DC area generalize to other locations with different land use, activity patterns, and data availability.

- **Benchmarks, metrics, and application-specific value need to evolve in parallel.** This study provides an initial characterization how well vendor data sources agree with one another and where possible, how closely they match reference datasets. These metrics and visuals serve as a starting point for quantifying vendor performance, but further work is needed to develop targets and benchmarks – a process that requires understanding agencies' intended use cases.

  For example, vendor products may be able to characterize spatial patterns – like identifying that Location A has twice the activity of Location B – well before they meet strict accuracy benchmarks for absolute volume. This ability to capture spatial patterns may be suitable for certain use cases (e.g., identifying high activity corridors, prioritizing projects, comparing exposure), but not others, reflecting the need to develop application-specific guidance.

Recommendations

In light of these findings, the validation team makes the following recommendations to agencies considering purchasing industry sourced bike/ped data products:

- **Make sure to understand product characteristics before purchasing.** Agencies are encouraged to review Part I of this report, which provides a summary and comparison of each vendor's product offerings. Because "bike/ped" data does not have a standard meaning, many diverse offers fall under the "bike/ped" label, even though some may not align with agency needs. For example, agencies should clarify whether they need zone or segment-level data and make sure that a candidate product supports the spatial and temporal granularity that is consistent with their intended applications.
- **Use industry sourced bike/ped data with caution.** Many of the industry-sourced bike/ped products included in this study appear promising and offer insights about bike or pedestrian activity in areas where no information currently exists.

  While this is exciting, *agencies should proceed with caution*, recognizing that many products are at an early product stage maturity and are evolving quickly. In particular, it is recommended that users:

  - *Prioritize applications that require relative measures of bike/ped activity, rather than absolute trip volumes.* Initially, it is likely that these datasets will be more useful for applications such as infrastructure justification, project prioritization, safety/risk analysis, and other use cases where relative spatial patterns are sufficient to make decisions. While it is expected that absolute trip volume accuracy will only improve over time – and may be reasonable for filling gaps where no other option exists -- the accuracy of reported trip volumes cannot be verified at the zone level (no 'ground truth'), and segment-level accuracy is not yet fully understood, particularly for pedestrian datasets (not studied in this analysis) and for different geographies and land use contexts that are different than the metro DC area.
  - *Avoid – or be careful – using industry-sourced bike/ped estimates for comparing bike/ped activity during different time periods*. It is difficult to know whether bike/ped metrics reported during two different time periods differ because of actual changes in bike/ped activity or other factors such as changes to a vendor's model/algorithm or

differences in underlying data supply. Some vendors may have reasonable ways to control some of these factors (e.g., re-generate prior year estimates with an updated model), but even so, the uncertainty associated with bike/ped trip estimates may obscure changes and trends over time.

Additionally, the validation team recommends that the Coalition take several steps to help industry-sourced bike/ped products continue to mature:

- **Encourage standardization of bike/ped product offerings.** In order for the industry to focus and converge on usable and replicable data sets, it would be beneficial for the Coalition to define a set of standardized offerings which enable productive DOT functions for planning and operation, and can be procured from multiple vendors, compared, and validated. This should be considered a priority in the next procurement of the Transportation Data Marketplace (TDM V4).
- **Support agencies with traditional bike/ped count programs and encourage standardization of count data.** Traditional, infrastructure-based counts are critical sources of calibration for industry-sourced bike/ped products, so having more and better-formatted agency count data would help vendors improve their products and also aid with ongoing validation efforts. TETC recently developed the Active Transportation Count Specification (ATCS) to serve as a common data model and exchange format for collecting and sharing bike, pedestrian, and other non-motorized counts. It is recommended that the Coalition include this specification in the upcoming TDM V4 procurement and encourage its adoption by supporting agencies with its implementation.
- **Develop metrics and application-specific guidance**. In the same way that guidance has been developed for motorized traffic volumes, it would be useful for Coalition to provide a framework for interpreting the quality of bike/ped data. This would likely involve outreach across Coalition agencies to understand expectations for accuracy by use case and in different contexts and may also include characterizing the error associated with traditional methods for obtaining annual average bike/ped volumes. For example, following the approach taken by FHWA for motorized volumes, this could involve characterizing the errors associated with factoring a 48-hr or 1-week bike counts to an AADT-equivalent value, and setting targets for accuracy/precision that would be at least as good as the 'status quo' approach.