STREETLIGHT InSight

Bike and Pedestrian Metrics Methodology, Data Sources, and Validation

Version 4.0

Updated September 2020

Table of Contents

Introduction

This technical document describes the data sources and methodology employed by StreetLight Data to develop our active transportation travel pattern Metrics -- bicycle and pedestrian, which are in our "Multimode" subscription – in addition to validation work to build confidence in the sources and data processing. This document is updated regularly – please check our website or reach out to your StreetLight contact to ensure you have the most update to date version of this document.

Locational Data Sources and Probe Technologies

Sources for Multimode Metrics

StreetLight Data's Multimode Metrics are currently derived from several types of data, predominantly:

1. General Location-Based Services (LBS) data
2. Mode-Tagged Location-Based Services (MT-LBS) data
3. Well-validated bicycle and pedestrian counts

As the mobile data supply landscape has evolved and matured over time, we have determined that these data sources are currently the best suited to analyze active mode transportation. In particular, we've evaluated and ruled out cellular tower data for active modes due to poor spatial precision and ping rate. Many Multimode trips are very short (<1 mile), so not knowing if a 300m variation is due to cellular tower accuracy fluctuations or an actual trip makes this an unsuitable source. Similarly, most navigation-GPS data comes from connected cars, making it useless for Multimode analytics. We found that using *only* mode-specific app data was not feasible due to low and biased sample size, with a bias towards recreational trips.

Table 1: Overview of Big Data supply options for transportation analytics. StreetLight recommends and uses a mix-and-match approach currently focused on navigation-GPS and LBS data types.

Type	Pros	Cons
Cellular Tower: **Derived from cellular tower "triangulation" and/or "multi-lateration" (100-2000m spatial precision)**	Large sample size - Most telecom providers have over 30M devices Ability to infer home/work locations	Very poor spatial precision Infrequent pings for some suppliers High cost Consumers must opt-out of data collection (vs. opt-in) Can't differentiate personal trips from commercial Poor coverage in rural areas No capture of short trips or ability to reliably infer active modes of transportation
In Vehicle Navigation-GPS: **From connected cars and trucks (3-5m spatial precision)**	Excellent spatial precision Very frequent pings Separates personal and commercial trips Opt-in for consumers	Usually lower sample size Difficulties inferring home/work (depending on supplier practices) No active modes, non-vehicular modes
Location-Based Services: **Mix of navigation-GPS, a-GPS, and sensor proximity data from apps that "foreground" and "background" with locational data collection (5-25m spatial precision)**	Very good spatial precision Frequent ping rate Superior ability to infer trip purpose and trip chains Ability to infer modes (walk/bike/transit/gig) Large and growing sample Opt-in for consumers	Less mature suppliers Variation in sample size and characteristics across suppliers requires more sophisticated data processing
Mode-Specific Location-Based Services Data: **Like LBS data, but from specific apps affiliated with modes**	Good spatial precision More certainty on mode, as users tag their mode (though we still find 20-30% points mis-tagged in these apps)	Much smaller sample size For active modes, sample is often skewed towards exercise/recreational biking and running

such as apps to log bike rides	Opt-in for consumers	
Ad-Network Derived Data: **When user sees an ad on their phone, their location is recorded by the ad network**	Large sample size of individuals	Few pings per month mean inference of travel patterns is not feasible

All of the data used for our Metrics requires significant cleaning, filtering, organizing, and algorithmic inference of trip structure and mode. These processes are described in the following sections.

GENERAL LOCATION-BASED SERVICES (LBS) DATA

LBS data can be processed into personal travel patterns at a comprehensive scale. Its high spatial precision and regular ping rate allow for capturing trips as well as activity patterns (i.e., home and work locations), trip purpose, and demographics, and enable mode inference to differentiate active from vehicular travel (more on this in sections below). This makes it an ideal alternative to data derived from cellular towers, which also has a large sample size but unfortunately lacks spatial precision and pings infrequently.

Our LBS data suppliers usually provide pieces of software called SDKs (software developer's kit) to makers of mobile apps to facilitate Location-Based Services. These smartphone apps include: couponing, dating, weather, tourism, productivity, locating nearby services (i.e.: restaurant/ bank/ gas station), and many more apps, all of which utilize their users' locations in the physical world in order to provide value to the users.

The apps collect anonymous user locations when they are operating in the foreground. In addition, these apps may collect anonymous user locations when operating in the background. This "background" data collection occurs when the device is moving. LBS software collects data with WiFi proximity, a-GPS (Assisted GPS) and several other technologies. In fact, locations may be collected when devices are without cell coverage or in airplane mode. Additionally, all the data that StreetLight uses has better than 20m spatial precision with an average ping interval of 200 secs/ping. Note – some of our LBS data suppliers build their own LBS collection code directly into the app, instead of using a common SDK, but the outcome is largely the same.

Mode-Specific Location-Based Services (MS-LBS) Data

In addition to general LBS data, we have special partnerships with apps that are dedicated to active mode transportation. For example, apps that are designed for people to track their daily walk, etc. These apps come with higher certainty that the trip made is in fact in the designated mode, so fewer points are mis-tagged. These apps have the same or better spatial precision as general LBS data, and often ping more frequently than general LBS data.

Well-Validated Bicycle and Pedestrian Counts

We use a set of well-validated bicycle and pedestrian counts to validate our Multimode Metrics. For example, a city may install a permanent loop counter in a bike lane, or use video-scanning technology to count pedestrians at an intersection. We also allow customers to "calibrate" (scale up and expand) our sample counts based on well-validated counters in that region. We attempt to use only permanent counts that we have deemed reasonable. We understand there is much debate around the exact accuracy of permanent (and temporary) count technology, which we do not address here. We do not use temporary counts in our default algorithms. However, customers may enter their own counts if they deem them accurate enough, no matter the collection method and duration, as a calibration zone. See streetlightdata.com/support for more on our calibration options. In the future, other data such as density, land use, weather, and more will be incorporated.

Additional Sources for Development of Algorithms

Additional sources of data were used during algorithmic development, though these sources do not directly provide data into the Metrics StreetLight provides to customers.

1. GPS-enabled travel diaries with detailed, user-entered, mode tags
2. A cohort of contract data trainers around the U.S., covering different urban forms and weather regimes, who walk or bike several miles a day and send their GPS logs and user-entered mode tags to StreetLight.
3. Traditional surveys about active mode behavior

The following section gives details on each data source used in the development of these StreetLight Metrics.

GPS-Enabled Travel Diaries With Detailed, User-Entered Mode Tags

Historically, many GPS-enhanced travel surveys have been performed. In some, users carried around GPS devices and then, via an app or online tool, tagged each trip with information such as the mode, trip end point, trip purpose, etc. While this data is private, as it describes personal activities, the data from several such surveys are housed at the Transportation Secure Data Center, hosted by NREL (https://www.nrel.gov/transportation/secure-transportation-data/). By applying to this group, and working fully within the databases' digital firewalls, we were able to use this data as calibration/ground truth data to develop machine-learning derived training algorithms that allow us to recognize walking or biking or (most difficult) transitioning between modes in our general LBS data. In addition, we enhanced the set with some similar published single-mode "pings" made public by bike share systems and academic papers.

Contract Multimode Data Trainers

To enhance our supplies of well-validated, user-tagged data, StreetLight has developed a network of data trainers across the U.S. This team walks, bikes, buses, train-rides, and drives around their towns (we've purposefully selected different places with different urban form across the U.S. and Canada) covering different types of roads, paths, and neighborhoods. They collect and "tag" their trips on their phone, similar to a GPS-enabled travel diary (see above) and send this data to StreetLight. We use this data to train and test (validate) our algorithms, and to see how well new iterations perform correctly interpreting the data from this team.

Traditional Surveys About Multimode

As we developed our algorithms, we used general characteristics of active transportation modes to help guide us. For example, we used distribution of bike-trip distance published by the national and regional household travel surveys. Say a local survey found that the average bike trip length in a city is four miles. We have two versions of our algorithm being tested. For version A, the average length of all bike trips in that city is five miles, for version B, the average is four miles. We will then favor version B.

Development and Implementation of our Mode-Tagging Algorithms and Metrics

Based on our deep knowledge of our own Big Data resources, and knowledge of travel patterns, we knew we needed a multi-faceted algorithm to measure active mode trips. For example, using average trip speed alone would not be adequate to characterize walk and bike trips, as a bike can go faster than a car in heavy congestion. Academic research on this topic concurs.[1,2]

However, we felt that much of the academic literature was not useful to us as it assumes that incoming data will be very refined (for example, pinging every second). We know that large, affordable samples of data are much messier and less granular. Hence, we needed to use an approach that accommodated the characteristics of our LBS data. We also felt that it was not reasonable to segment our data into trips, and then infer mode. That's because different modes have different ways to end the trip. The end of a walk trip, as someone enters a building and walks at the same speed to their desk looks very different than the end of a car trip, when someone goes from 35-mph to 0-mph very quickly, and then walks a few feet to their front door.

Therefore, we decided that the best way to approach mode inference was to:

[1] Dabiri and Heaslip. "Inferring transportation modes from GPS trajectories using a convolutional neural network." Transportation Research Part C 86 (2018) 360-371.
[2] Wu, Yang, and Jing. "Travel mode detection based on GPS raw data collected by smartphones: a systematic review of the existing methodologies." Information 2016, 7, 67; doi:10.3390/info7040067.

1) Assign mode probabilities to every "ping." We chose to model a device at rest, when it continues to ping, as a "stationary" mode.
2) Group pings into trips, with mode-adaptive criteria for breaking trips.
3) Assign the most probable mode to each trip, given ping probabilities combined with overall characteristics of the entire trip.

It should also be noted that we've continued to evaluate and update our algorithm in order to provide the highest quality Metrics possible. Our latest iteration applies to months from January 2019 onward. When updating the algorithm for the 2019 and 2020 Multimode Metrics, we made the following enhancements:

1. We incorporated vehicular data from Navigation-GPS devices to help with mode classification (to help differentiate fast bike from car trips, for example)
2. We considered additional geospatial features (outlined in Step 1 below) to assist with mode differentiation

Step 1 – Assign Mode Probabilities to Every Ping

For this step we used machine learning techniques to assign mode probabilities to every ping. After considering several techniques, we decided that a random forest (RF) classifier would yield the most accurate results while still yielding an algorithm that can be efficiently implemented in our product at scale.

TRAINING DATA

The first part of a successful machine learning project is to develop a very clean set of "training" or calibration data. This data must contain pings (single points with location and time) that comprise many trips, to which the user has confirmed that they were walking at this time, biking at this time, etc. To do this, we mainly relied on NREL's Transportation Secure Data Center.

Some of the training data came from surveys that had volunteers wear a GPS device to record various trips and tag them with the travel mode as they went about their daily activities. The GPS devices would log the user's location (as latitude/longitude) and their timestamp. More detail can be found at the NREL TSDC website.[3] We used the following surveys whose data is housed in NREL's TSDC:

- California Dept. of Transportation – 2010-2012 California Household Travel Survey
- Atlanta Regional Council – 2011 Regional Travel Survey
- Mid-America Regional Council – 2004 Kansas City Regional Household Travel Survey
- Mid-Region Council of Governments – 2013 Mid-Region Travel Survey
- Regional Transportation Commission of Southern Nevada – 2014 Southern Nevada Household Travel Survey

[3] https://www.nrel.gov/transportation/secure-transportation-data/

In addition, we used smaller sets of similar data from bike shares and data supplied by academic papers published on this topic. Much manual quality assurance was done on all training data, as it contains many inaccurately tagged points. We understand that some of these surveys are from years in the past. However, we need the data to accurately represent the cadences of a walking person or cyclist, etc. This, we presume, has not changed in the past few years. We do not need the data to be statistically representative across space or time (i.e., we do not need a lot of trips), however we do need many examples of each mode of travel.

In our 2019/2020 algorithm, we also incorporated vehicular data from Navigation-GPS devices. This allowed us to consider a wider breadth of car trips, including more slow-moving vehicles, ultimately helping to reduce mis-classification with other modes.

Finally, the data collected by our “StreetLight-specific” team of data trainers across the U.S. (as described above) gives us an even richer and more varied set of training data.

We had to alter the training data so that it looked as similar as possible to our core LBS data. In particular, most of the training data pinged far more frequently than our LBS data; therefore, we impoverished the training data by removing points.

Model Features and Model Training

We experimented with over 50 features. Some are “interior” to the pinging geospatial data including time, distance, speed, acceleration, jerk, circuity, and angular velocity for each ping as well as for its preceding and subsequent pings, day of week, hour of day, etc. Others are “exterior” or “contextual” including road classification, weather, and density of commercial activity nearby. In our 2019/2020 model, we considered additional geospatial features. This included the presence of bike and bus lanes, proximity to parks, as well as road network density. Incorporating more geospatial features into the model allowed us to capture more bike and walk trips on trails, and also better differentiate between modes in urban areas.

Features can be thought of as attributes or explanatory variables in a model. In the end, we used a subset of the features which were most impactful in the training of the random forest algorithm. Also, we did not want to use too many features to avoid over-fitting, and instead allowed the model to be more adaptable.

To train the machine learning model we used a classic 80-20 split, training the model on 80% of the training data and testing it against the remaining 20% which we’d held back from training. We used a technique called “bagging” to try this iteratively with 100 decision trees in the random forest to arrive at an “out-of-bag score” which we could use as a measure of the quality of predictions from the model. Furthermore, we also monitored and made changes to improve the precision/recall scores and F-scores for each mode (car, bike, walk) which the model was supposed to be able to predict.

Our model is constantly improving. Some versions of the model score higher but use much deeper trees (which have a higher computational cost to run). Our selected version balances high scores with maintaining cost and efficiency for our customers at scale.

The modes predicted by the RF classifier are only reflective of the immediate vicinity of the ping. There can be a lot of noise and variation from ping to ping in the span of an entire trip. A vehicle driving down a congested road may be slower than a person walking down the road. A bicyclist riding a road bicycle may at times travel faster than a car.

Hence, we took a probabilistic approach to consume the results of the classifier. When we run our algorithm, pings are not assigned a single mode. Instead, they are assigned a mode probability distribution (0 – 1). For example, Table 2 lists a time sequence of pings where a person transitions from driving a car to walking for ~10 minutes.

Table 2: Time sequence of pings for a mode transition. This demonstrates the probabilistic approach used to classify individual points.

Individual Mode Probabilities					
Ping Timestamp	Bike	Car	Stationary	Walk	Prediction
5/28/2019 10:51	0.036	0.939	0.007	0.018	Car
5/28/2019 10:53	0.087	0.872	0.003	0.038	Car
5/28/2019 10:55	0.109	0.752	0.003	0.136	Car
5/28/2019 10:58	0.261	0.345	0.137	0.257	Car
5/28/2019 11:00	0.215	0.320	0.000	0.465	Walk
5/28/2019 11:02	0.162	0.165	0.005	0.668	Walk
5/28/2019 11:04	0.204	0.100	0.003	0.693	Walk
5/28/2019 11:07	0.082	0.037	0.008	0.873	Walk
5/28/2019 11:09	0.049	0.036	0.002	0.913	Walk
5/28/2019 11:11	0.048	0.063	0.001	0.888	Walk

Steps 2 and 3 – Group Pings Into Trips, and Assign Modes to Trips

We have a “sequence-of-linked-pings” approach where the linkage is based on identifying the same mode or a similar mode of the pings ordered by their timestamp. The “sequence-of-linked-pings” grows into what is eventually determined to be a trip.

In a nutshell, we endeavor to first determine the trip boundary where one trip ends and another begins, then decipher the ending trip’s probabilistic travel modes by calculating it from the individual pings in the trip. Given the probabilistic nature of our approach, we consider a primary and a secondary, etc., mode for each ping with an associated normalized probability. Similarly, we derive a primary, secondary, etc., mode for each trip along with an associated normalized probability for each mode.

We read the pings ordered by their timestamp, predict the travel mode of each ping using the machine-learning model-based classifier (as described in step 1) and send the stream of mode-tagged pings to the process that creates trips. When we encounter enough consecutive pings with a different enough set of modes, we conclude that the mode has changed. We then end the current trip and start the next trip with the new mode.

A note on “stationary” trips. As mentioned in step 1, we have chosen to model a person at rest (i.e., a person who is not moving but stationary) to have a "stationary" travel mode. This helps model the process of a person in motion coming to rest (because his/her real-world trip ended) as a mode change by the system since the travel mode of the person changed from, say, "car" to "stationary" travel mode. In our 2019 and 2020 trip creation process, we made additional enhancements to our stationary strategy, contributing to more effective trip trimming, and improving the quality of trip starts and ends.

Similarly, when a new trip is started and the person starts moving again, the travel mode changes from "stationary" to, say, "bike." This technique makes the starting and ending of motion harmonized as a simple mode change and indistinguishable from, say, a person walking to a car and driving (when it would be a mode change from “walk” to “car”).

As a final check, we verify the correctness and feasibility of each trip for the mode assigned to it. For instance, if the average speed of a bike trip is higher than a certain pre-defined "impossible" speed, it is eliminated. We do not rule out bike trips if they happen on roads with no bike facilities, as this happens often. Similarly, we do not rule out pedestrian trips if there are no footpaths or curbs. If we find a trip that appears to be missing its end or beginning – for example, a car trip that stops in the middle of the highway – we eliminate it. If we find a trip that appears to have erroneous data – for example goes from San Diego to Africa and back in four seconds – we eliminate it.

Step 4 – Locking

A trip from an LBS device is a series of connected pings. If the traveler turns a corner but the device is only pinging every 10 seconds, then that intersection might be “missed” when all the device’s pings are connected to form a trip. For vehicular modes, StreetLight utilizes road network information from OpenStreetMap (OSM) including speed limits and directionality to “lock” the trip to the road network. This “locking” process ensures that the complete route of the vehicle is represented, even though discrepancies in ping frequency may occur.

However, for Multimode Metrics, locking is a more ambiguous task. Pedestrians and bikes do not always follow the rules around speed and directionality. Thus, we allow pedestrian and bike trips to be locked to the best OSM segment (including the OSM bike-only and ped-only networks) no matter the presence of a sidewalk, or bike permissions, etc. Bike trips follow one-way rules of the road, but pedestrian trips do not have to.

Step 5 – Contextualization: Demographics and Trip Purpose Assignment

If a device that creates LBS data regularly pings on a block that contains residences, and those pings often occur overnight, there's a high probability that that the device's owner lives on that block. This allows us to associate a "likely home location" to that device. By using frequency and land-use context, we avoid assigning a device to "live" at a place where the owner may work a night shift (an airport for example) or goes on a vacation just for a few days a month. We reset each device's home location each month to accommodate people who move residences.

The trips that have the home location as a trip start or trip end can be called home-based. In addition, we can append distribution of income and other demographics for residents of the matching Census blocks (or block groups, for some data) to that device. That device can then "carry" that distribution everywhere else it goes. Our demographic data sources for the U.S. are the 2010 Census and American Community Surveys. In Canada, our source is Manifold Data. In addition, the home assignment supports our normalization process, discussed below.

To assign workplaces, we look for where a device most frequently spends daytime hours. We do not use land use for this assignment, since work occurs in all land uses. Note that we may allow people to work from the same place as their home. We will typically not capture people who work from home in our Home-Based Work trip purpose, since they will not have a "trip" between home and work. We do allow people to "work" at places like college campuses, thus students going to school may be classified as working at the school. We will also miss people who work at irregular locations, like plumbers.

Step 6 – Store Clean Data in Secure Data Repository

After being made into patterns, checked for quality assurance, normalized, and contextualized, the data is stored in a proprietary format. This enables responses to queries via the StreetLight InSight® platform in an extremely efficient manner. By the time the data reaches this step, it takes up less than 5% of the initial space of the data before the extract, transform and load (ETL) process. However, no information has been lost, and contextual richness has been added.

Step 7 – Quality Assurance

Quality Assurance is done at multiple levels:

- Classification of individual pings: For the testing of the classifier we used typical techniques used in testing machine-learning algorithms. We trained the model on 80% of the training data and tested it against the 20% which we had held back and not exposed to the model. This is discussed in the sections above.

- Unit-level testing for trips: For testing the creation and breaking of the trips we hand-picked 100+ unit test trips against which we verified the trip boundaries, the overall trip travel mode, as well as the mode of the individual pings as determined by the system.
- System-level testing for trips: We performed a number of tests on the overall trips generated in each month of data, including visual checks, statistical checks, spatial and spatio-temporal checks, and real-world data comparisons.

In addition, bike trips longer than 50 miles (before locking) and pedestrian trips longer than 10 miles (before locking) are considered erroneous and eliminated from the sample.

Some data from this QA process is available in validation papers, available on our website.

Step 8 – Normalization and Expansion

For bicycle and pedestrian trips, StreetLight uses a set of permanent bike and pedestrian counters on trails and bike paths to measure the change in trip activity each month. Then we compare this ratio to the ratio of trips at the location, and normalize appropriately. For example, if the permanent counter says there are 400 bike trips a day, and we sense 40, the expansion ratio will be 10. Thus, the StreetLight Index for Muiltimode data is normalized to adjust for change in our sample size. We performed a thorough review of monthly changes in our sample and ensured that our normalization process allowed for comparisons across time.

Due to varied normalization ratios across road types and geographies, the StreetLight Index for Muiltimode data is not yet “expanded” to estimate the actual flow of travel. Downloads are available with all Metrics shown in sample counts, and the StreetLight Index.

For customers who want to expand to estimated bike counts, we have a tool called “Calibration” which allows the user to enter bike or pedestrian count data from an external source, such as a sensor, and then automatically integrate that data with our LBS-derived data to get an estimated count for any nearby pass-through zone, O-D pair, or any other StreetLight Metric. Some customers do additional expansion or factoring on their own after downloading StreetLight Metrics and working with the sample counts provided for each zone.

Step 9 – Aggregate in Response to Queries

Whenever a user runs a query via StreetLight InSight®, our platform automatically pulls the relevant trips from the data repository and aggregates the results. For example, if a user wants to know the share of trips from origin zone A to destination zone B vs. destination zone C from September 2019, they specify these parameters in StreetLight InSight. Trips that originated in origin zone A and ended in either destination zone B or destination C during September 2019 will be pulled from the data repositories, aggregated appropriately, and organized into the desired Metrics. Several assumptions are made while aggregating, and can be found in the FAQ section of our support center (www.streetlightdata.com/support).

For example, when grouping trips into daypart bins (different time periods during a day) for an origin-destination analysis, the time that the trip begins is used to define the time-period bin for the trip. But when looking at dayparts for an origin-destination analysis with select link, the time that the trip traverses the select link is used. For active mode trips, only trips that are highly probably bike or pedestrian (as selected by the app user) will be queried. When setting up an analysis in StreetLight InSight®, users are able to specify the desired mode when selecting the "Mode of Travel" at a first step in the "Create Analysis" process. See our support center for guides on how to set up a bicycle or pedestrian analysis.

StreetLight is constantly improving its Multimodal Metric. Please tell us what you find! We will share future improvements in methodology and validation in updates of this white paper, so check our website regularly for new information.

Bicycle Validation

Understanding existing bicycle and pedestrian behavior is crucial to active transportation and planning efforts. StreetLight has developed algorithms and machine learning techniques that utilize Location-Based Services data in order to identify bicycle and pedestrian trips across the United States and Canada. This validation section focuses on comparisons between StreetLight Multimode Metrics, published travel survey metrics, and permanent counter locations across the states of California and Pennsylvania, as well as Ottawa, Ontario in Canada.

We've taken two approaches to validate our Multimode Metrics. First, we compare our aggregated trip characteristics to information published by household travel surveys in order to evaluate whether our Metrics are within the general range of expected trip characteristics for each mode. Second, we compare our trip volumes for particular locations (roads and designated bike paths) to published bicycle and pedestrian counts from permanent counters.

We attempt to use only permanent counts that we have deemed reasonably accurate. We understand there is debate around the exact accuracy of permanent count technology (and we found many erroneous points within the published data). We do not use temporary or expanded counts for validation, as they have additional sources of error.

Our goal in this validation section is to demonstrate that our Multimode Metrics can be used in place of surveys or temporary counts.

Comparing Our Bicycle Results to Travel Surveys – NHTS and CHTS

We analyzed our aggregated bicycle trip attributes and compared them to the National Household Travel Survey 2017 (NHTS) and the California Household Travel Survey (CHTS) 2012-2013.

To create a comparison metric set, StreetLight Data analyzed roughly 126.7 million bicycle trips across the continental U.S. that occurred from January 2019 through May 2020 for an average of 7.5 million trips per month (for more on how we infer and create bike trips, please see the separate methodology section). NHTS surveyed ~129,000 households across the county and analyzed 8,000 bike trips while CHTS surveyed 42,431 households with fewer than 5,900 bike trips.

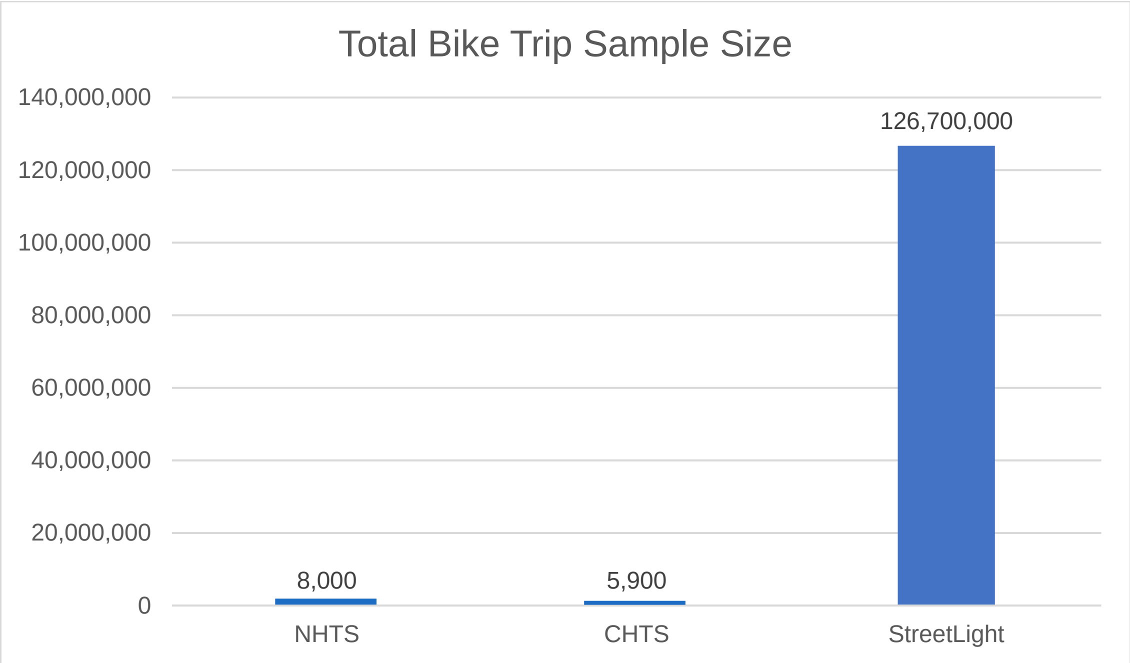

Figure 1: Chart comparing NHTS and CHTS total bike trips sample sizes to StreetLight's bike trip sample size.

We compared key average characteristics of biking from our data set to the household surveys. We do not set an "exact match" as the goal. All three data sets are samples, and all three thus have different strengths and sources of error. Where discrepancies occurred, we believe they are explainable by known differences in collection methods. The following tables show StreetLight average bicycle trip attribute Metrics relative to the published NHTS and CHTS numbers.

Table 3: Comparison of average trip length from different sources (miles).

Mode	NHTS	CHTS	StreetLight - National
Bicycle (all)	2.38	1.50	2.55

As shown in Table 3, we found the StreetLight trip distance averages were very close to the NHTS, and slightly longer than the CHTS. Figure 2 shows the distribution of bike trips by length. StreetLight's distribution trends match very closely with survey, though we do capture more long-distance bike trips than the survey, with a less defined peak at one mile compared to NHTS results.

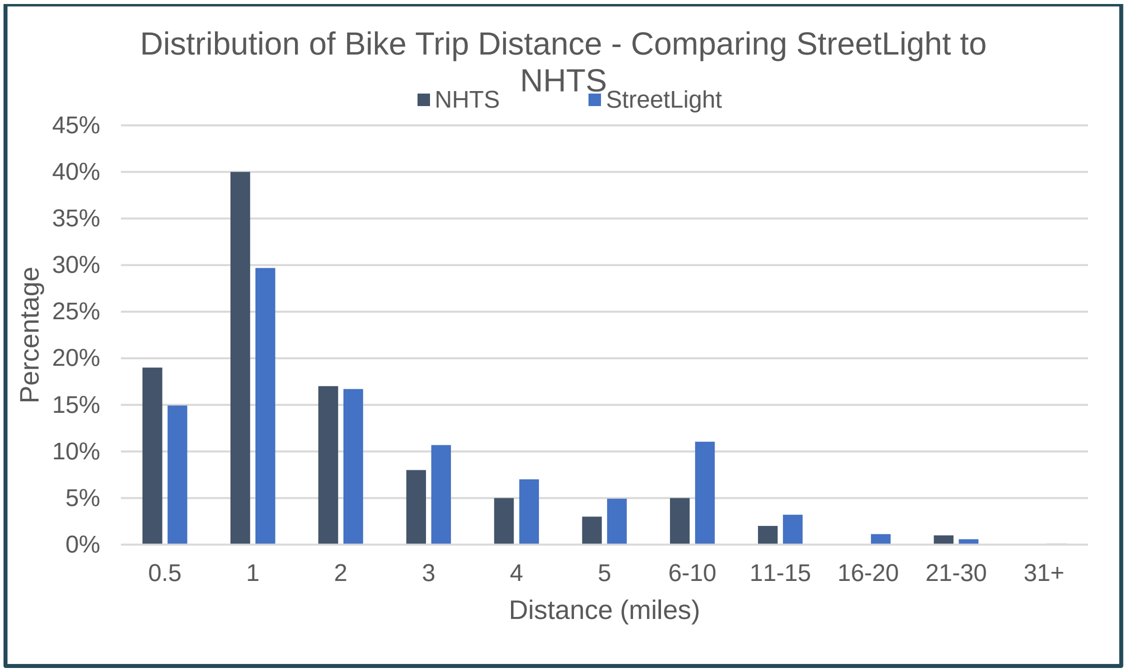

Figure 2: Histogram of bike trip length distribution for StreetLight and the NHTS.

It is not possible to determine which is "more correct" as all sources come from a sample. However, there are some inherent differences between surveys and passive Big Data that should be considered when evaluating the strengths and weaknesses of both sources, specifically:

1. **Trip-breaking criteria:** StreetLight "breaks" bike trips using a layered approach with multiple conditions. At a high level, if the machine learning algorithm sees two consecutive points beyond the current trip boundary whose combined probability of biking is < 50% then we conclude the person may have "changed" modes, which includes coming to a halt since we model "standing still" as a mode in itself. Thus, a bike trip that involves biking along a trail, stopping for lunch at a scenic overview, then continuing the ride would show up as one long trip in a survey, and two shorter trips in StreetLight.
2. **GPS-assist:** The CHTS used GPS assist to help measure this trip-length bias and found that shorter trips were far more likely to be under-reported than long ones.[4] They then adjusted results to help correct for this bias. They may have over adjusted.

Next, we compared trip time and average trip speed to the survey results. Overall, we observed that the StreetLight average bicycle trip duration is similar, though slightly longer, than NHTS and CHTS metrics, as shown in Table 4. Some of this is explained by the difference in average trip length, reported above.

[4] Zmud, Johanna. "Identifying the Correlates of Trip Misreporting - Results from the California Statewide Household Travel Survey GPS Study." 10th International Conference on Travel Behavior. August 2003.

Table 4: Comparison of average trip duration (minutes).

Mode	NHTS	CHTS	StreetLight Trips - National
Bike	21.92	18.20	24.50

Table 5: Comparison of average trip speed (mph).

Mode	NHTS - Imputed	CHTS - Imputed	StreetLight Trips – National
Bike (avg. distance / avg. duration)	6.51	4.95	7.20

We noticed a slight increase in speed from prior versions of our algorithm, likely due to the fact that we now more accurately classify high-speed bike trips on trails. Overall speed comparisons align well with NHTS.

Comparing Our Results to Permanent Bike Counts

To validate data at a more granular level we compared bike volume counts for months between 2019 and 2020 at specific roads to permanent bike counters across cities nationwide.

Bike Counter Source Review

For this validation, we aimed to analyze permanent bike counters with data available through Spring 2020 in order to be able to validate as many of StreetLight's available months as possible. Since many sources for permanent count data are published annually, this left us with a limited set of potential locations for comparison. We needed counters that published their data on a daily or monthly basis in order to meet our strict criteria for temporal coverage. We were able to find high-quality sources spanning various U.S. and Canadian cities.

The first city we chose was San Francisco because of the high number and diverse locations of permanent continuous bike counters, and the frequency of their data updates. This data is made available by the San Francisco Municipal Transportation Agency (SFMTA).[5] The SFMTA dataset has 27 bike counters that collected data across our analysis months. Only 11 of those 27 permanent counters recorded data over all 17 months we were looking to analyze (January 2019 – May 2020). SFMTA's viable counters included 25 that were on urban roads, many of those with dedicated bike lanes, and two that were on separated trails or paths, independent from the vehicular road network.

[5] https://www.sfmta.com/bicycle-ridership-data-1

Philadelphia contributed additional reliable counts, made available by the Delaware Valley Regional Planning Commission (DVRPC)[6]. DVRPC maintains a number of permanent counters across the city and in some surrounding suburbs. In total, there were 14 permanent counters that collected data in the analysis months. Of those 14 counters, only four provided data consistently across the 17-month period. All 14 of the permanent counters were located on trails, usually bike and pedestrian paths in parks in Philadelphia and the greater Delaware Valley.

We also sourced an additional permanent counter in Canada on a bike and pedestrian bridge (Adawe Crossing) in Ottawa, Ontario. This data was made available by the City of Ottawa.[7]

The SFMTA, DVRPC, and City of Ottawa data sources reported bike counts hourly over the 17-month period. Thus, we were able to also aggregate daily and monthly totals from the datasets. We performed additional filtering on the monthly counts to ensure locations and months with wild swings were not included in our comparisons.

In evaluating the permanent counters for comparison, it's important to consider the range of average daily trip counts recorded across the locations. The fundamental difference in average daily trips is tied to whether the counter is located on a road or a trail. Roads, most of which are in urban locations, record a higher daily average than trails, as seen in Figures 3 and 4 below. We will discuss these differences in more detail later on in this validation.

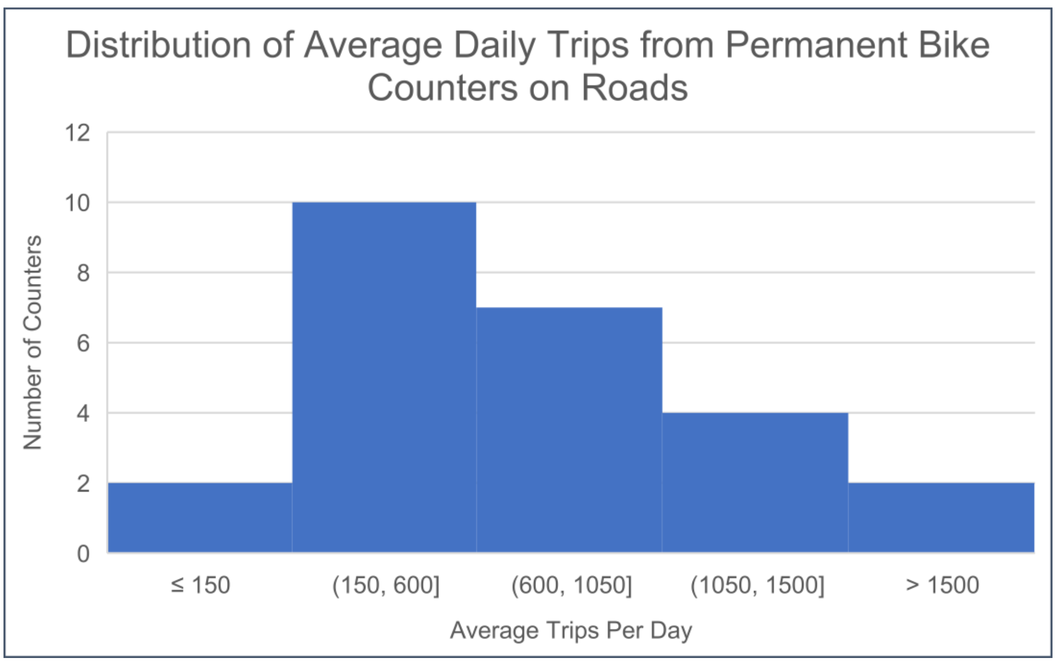

Figure 3: Histogram of selected permanent counters' average trips per day on roads. Most locations record less than 1,000 trips per day. Counters are located in San Francisco.

[6] https://www.dvrpc.org/webmaps/pedbikecounts/

[7] https://open.ottawa.ca/datasets/bicycle-trip-counters

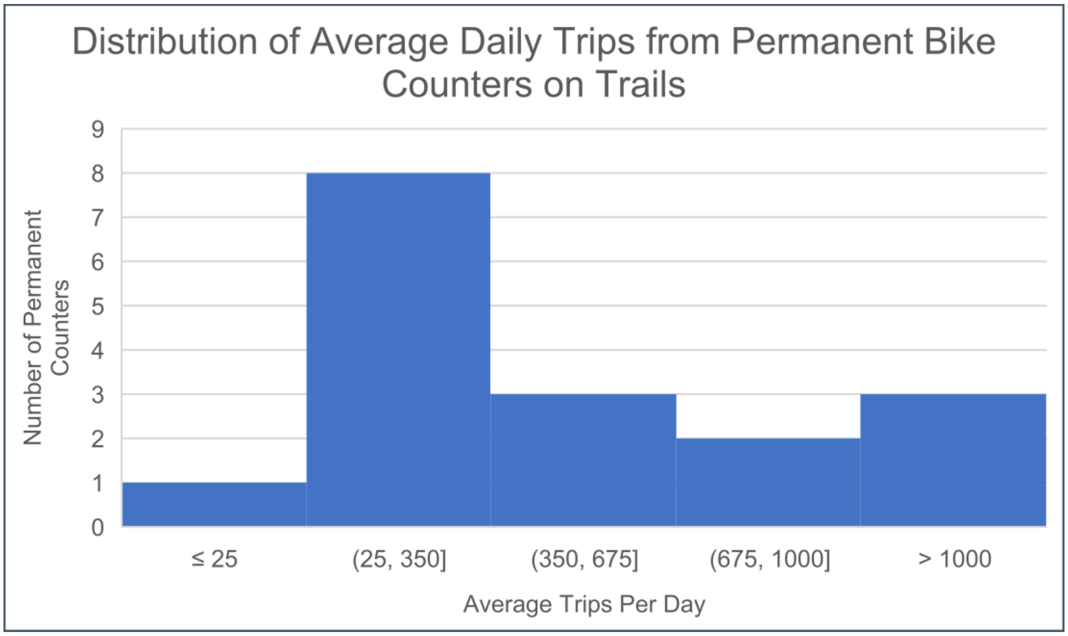

Figure 4: Histogram of selected permanent counters' average trips per day on trails. Most locations record less than 350 trips per day. Counters are located in San Francisco, Philadelphia and Ottawa.

INITIAL COMPARISON TO STREETLIGHT

After the source data was cleaned, we compared the monthly counts at these counters with the bicycle counts on the matching segments obtained from the Streetlight InSight® platform. Our goal was to determine the correlation between the two sets of counts. Since we were working with a limited number of permanent counters, we wanted to include as many months as possible in our comparison, and not necessarily fully exclude counters with missing months, or highly variable months (indicating some counter error). This means that if a counter had months missing, or highly variable outlier months between January 2019 and May 2020, its viable months could still be used for comparison. Therefore, the aggregated comparisons performed for some counters may be the sum of all 17 months, while others may be a sum of a smaller subset of months.

Before diving into deeper validation, we wanted to highlight some potential widely documented sources of error with permanent bike counters:

1. Counters only record trips that pass directly over the counter area. Thus, if bicyclists are traveling outside the bike lanes, they will be missed.
2. Some bike counting technology will under sample bikes made of particular materials, or confuse bikes with cars or scooters.[8]
3. Permanent counters can degrade over time with broken hoses or loops which can also be covered during road repair projects.

[8] https://www.fhwa.dot.gov/policyinformation/travel_monitoring/pubs/pedbikedata.cfm#sect7_3

4. Like any data collection, errors may occur in the transferring of data and posting of it to the online database.

The comparison of counters to StreetLight Metrics has an additional source of potential error: trip "locking" or assigning a bike trip to a particular route, based on waypoints measured during the trip. This is a bigger source of error for bikes than for cars, as bicyclists can take advantage of different flow lanes and bike-only facilities, making them harder to lock. They also may not take the fastest or most logical route because of hills or safety.

In order to get a sense of our general performance across counter locations, it's important to consider the rates of penetration recorded at each counter. Penetration rates are calculated as the sum of the monthly StreetLight sample divided by sum of the monthly counter sample. A relatively consistent penetration rate is key to consistent results in a validation. Penetration rates are also key to helping us understand which counters and months are outliers, indicating some other source of error may be at play.

In analyzing our selected permanent counters, we determined that the distinction between trails and roads was important in order to understand how our data performs across geographies and types of facilities. Generally, we see that roads have higher penetration rates than trails. While our available road counters were concentrated in San Francisco, our trail counters were dispersed across San Francisco, Philadelphia, and Ottawa, suggesting these patterns in our penetration rates were consistent across regions.

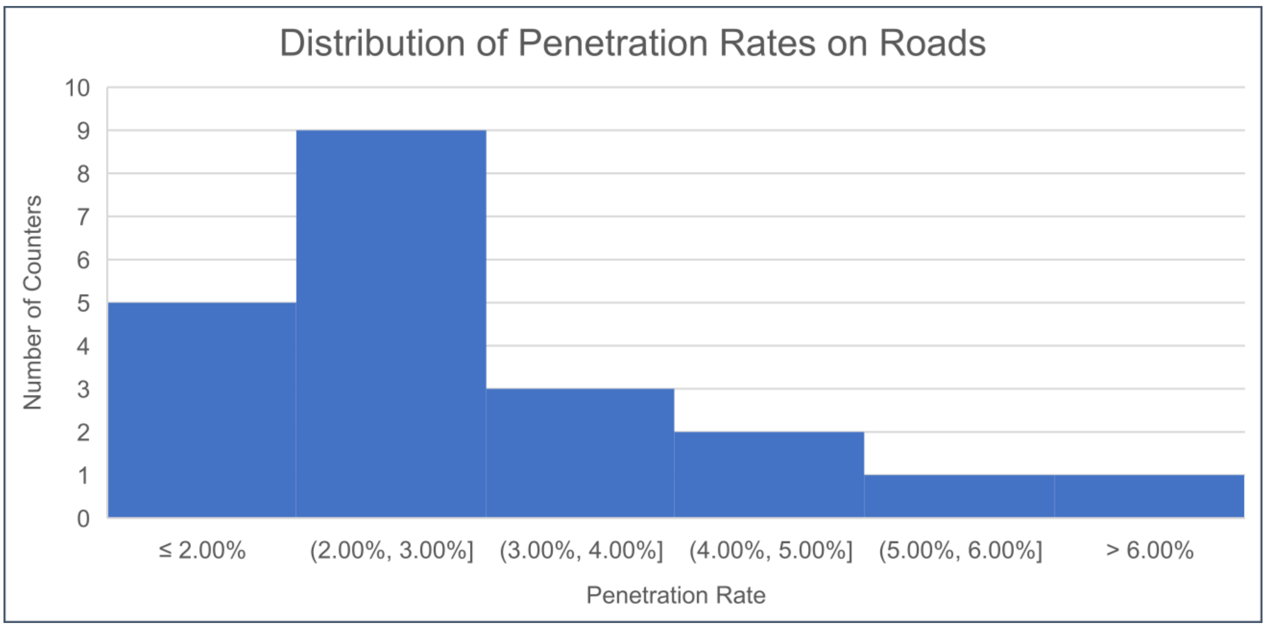

Figure 5: Distribution of average penetration rates on roads. Most locations see penetration rates between 2 and 3%. Counters are located in San Francisco.

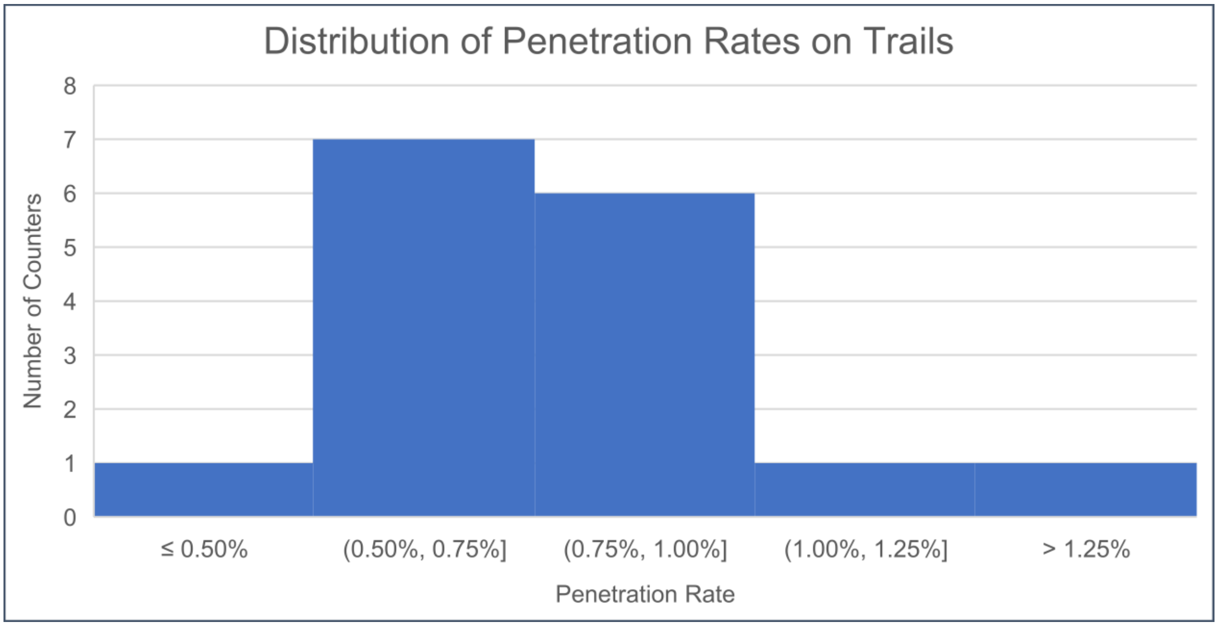

Figure 6: Distribution of average penetration rates on trails. Most locations see penetration rates between 0.50 and 1%. Counters are located in San Francisco, Philadelphia, and Ottawa.

It should be noted that earlier versions of our machine learning algorithm sometimes had trouble capturing fast-moving bicycles on trails. With the algorithm updates implemented for our 2019/ 2020 Multimode Metrics, we were able to classify more bicycle counts on trails by using geospatial features such as considering proximity to parks, trails, or main roads as well as increased amounts of training data. We will continue to improve this approach, which may lead to a closer match between penetration rates. In addition, in the future we will explore possibilities resulting from different phone usage behavior while on trails (leaving phones locked in the car, turning off geolocation features to conserve batteries, for example).

TIME TRENDS

In validating our results, we first decided to evaluate hourly distribution across average weekends and weekdays. For the day type and day part comparison, we first looked at the sum of trips across hours for all locations to see if the distribution was consistent. Figures 7 and 8 show the hourly distribution of trips in San Francisco and Philadelphia.

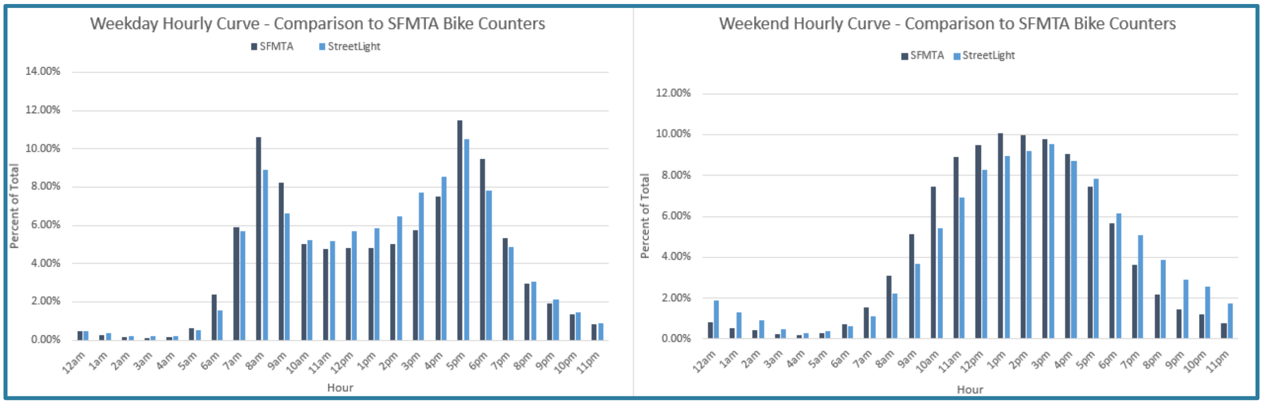

Figure 7: Average hourly curves on weekdays and weekends, comparing SFMTA's permanent counters to StreetLight Metrics at the same locations between Jan 2019 and May 2020.

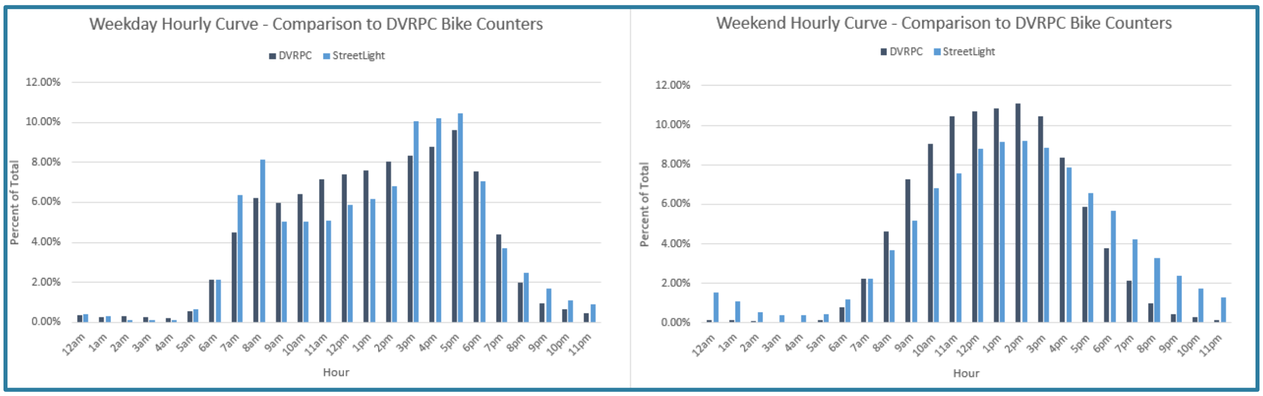

Figure 8: Average hourly curves on weekdays and weekends, comparing DVRPC's permanent counters to StreetLight Metrics at the same locations between Jan 2019 and May 2020.

What's notable is the different shapes of the curves for weekdays in San Francisco relative to weekdays in Philadelphia. San Francisco zones, which are primarily on roads, follow a clear commuter pattern on weekdays, with sharp morning and afternoon peaks. Philadelphia, with zones primarily on trails, shows a slight morning peak, but a steady increase in trips throughout the day. Our ability to mirror these unique hourly trends confirms our confidence in comparing weekdays to weekends, as well as comparing day parts to one another.

In general, but on weekends in particular, StreetLight under-samples compared to SFMTA and DVRPC during the mornings, compared to early afternoons. This requires further investigation and potential algorithm adjustments, but may also be accounted for in factoring and normalization algorithms in the future. Overall, the individual count locations showed similar trends compared to the aggregate.

In addition to evaluating hourly trends, we also wanted to confirm that our normalization process enabled users to capture monthly trends. Our normalization process, which uses monthly scaling values derived from comparisons to permanent counters, should ideally capture the seasonal or monthly curves recorded by the permanent counters. Since we are using a single scaling value (reference the normalization step in our earlier methodology section), we are not expecting to produce estimated counts that perfectly mirror the monthly permanent counts. However, we are looking to capture the variation across time so that months can be directly compared to one another using the StreetLight Index. So as not to bias the results, none of the locations listed below in Figure 9 were utilized to create the scaling values for monthly normalization.

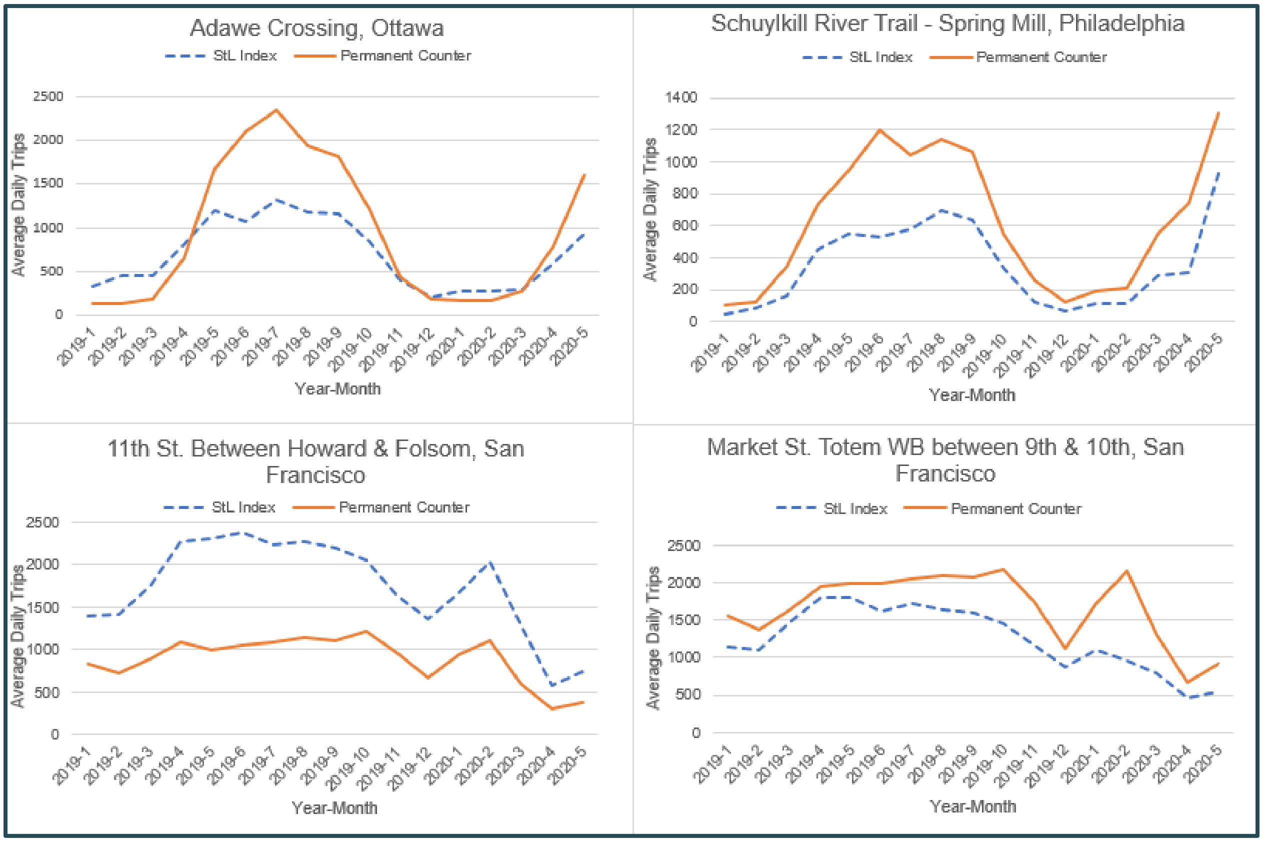

Figure 9: Average monthly curves, comparing permanent counter locations in San Francisco, Philadelphia, and Ottawa to the StreetLight Index between Jan 2019 and May 2020. For each chart, the year and month are represented on the x-axis, while the average daily trips are represented on the y-axis.

The charts above indicate the StreetLight Index is able to capture seasonal and monthly curves reported by the permanent counters. We're pleased that the StreetLight Index is able to capture seasonal peaks, as well as changes during 2020 months impacted by COVID-19 (March through May 2020). A notable trend is the way trails and roads indicate different patterns during COVID-19. Bike activity on roads heavily used for commuting (the two San Francisco locations) see steep declines in average daily trips, while trails, which likely have more recreational trips see standard seasonal increases. In the case of the Schuylkill River Trail - Spring Mill in Philadelphia, we see an even steeper increase indicating higher usage during stay-at-home mandates.

COMBINED CORRELATION RESULTS

To evaluate correlation between counter locations, we combined the SFMTA, DVRPC, and City of Ottawa results. Rather than evaluating the correlation by city, we decided to evaluate the correlation by facility type, splitting out roads from trails. As we saw with our earlier evaluation of daily counts from permanent counters, as well as our rates of penetration, the road and trail distinctions are key to understanding the performance of our bike Metrics.

First, we evaluated the correlation between StreetLight counts and the sum of average daily trips, recorded by month, for each permanent counter. For counter and month combinations with penetration rates outside of our expectations, we removed those months as outliers, and summed the remaining non-outlier months for the given counter. Based on our observed spread of monthly penetration rates, we considered outliers to be months with penetration rates above 7% for roads, and months with penetration rates above 7% or below 0.25% for trails. Many trails had very low monthly counts, thus we want to ensure a sufficient monthly sample. This allowed us to handle cases where error may have been introduced into a single month, without impacting all remaining months.

We evaluated 2019 months individually, and then evaluated the combined results for 2019 and 2020 as well. For roads, the results showed high correlation, with an R^2 value of 0.80 for 2019 months only, and 0.77 for 2019 and 2020 months combined. We believe this difference can be explained by the fact that 2020 months on roads were more variable, due to the impact of COVID-19 which led to decreased trip samples on commuter-heavy facilities.

This means that StreetLight's sample alone "explains" 80% of the variation in the permanent counters on roads in 2019. No factoring or adjustment has been done to these results. The FHWA recommends that even permanent counters be factored based on location type, sensor type, season/weather, and more to compensate for inaccuracies.[9]

[9] https://www.fhwa.dot.gov/policyinformation/travel_monitoring/pubs/pedbikedata.cfm#sect7_3

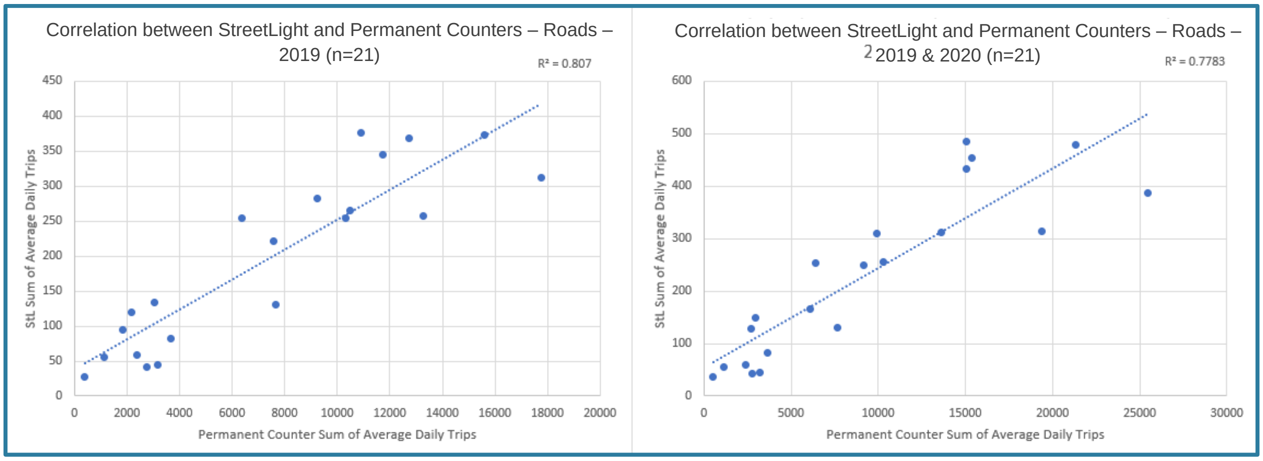

Figure 10: Correlation between average daily bike trips reported by permanent counters and average daily bike trips reported by StreetLight on roads. The first chart (left) indicates correlation for 2019 months with an R^2 value of 0.80. The second chart (right) indicates correlation for 2019 and 2020 months combined with an R^2 value of 0.77. Counter locations are in San Francisco.

For trails, we also see high correlation, with an R^2 value of 0.90 for 2019 months only, and 0.90 for 2019 and 2020 months combined. As described earlier, trails saw less dramatic changes in traffic during months impacted by COVID-19, and even increases in daily trips at some locations. As a result, we see similar quality results when 2020 months are included in the analysis.

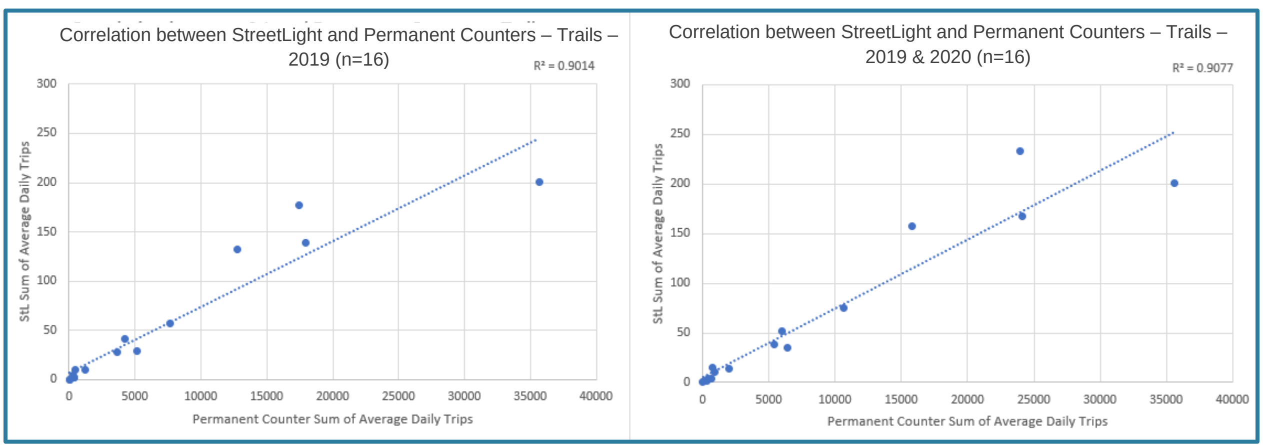

Figure 11: Correlation between average daily bike trips reported by permanent counters and average daily bike trips reported by StreetLight on trails. The first chart (left) indicates correlation for 2019 months with an R^2 value of 0.90. The second chart (right) indicates correlation for 2019 and 2020 months combined with an R^2 value of 0.90. Counter locations are in San Francisco, Philadelphia, and Ottawa.

We know many customers are interested in evaluating commuter patterns, so we wanted to ensure we had strong results when evaluating trips on average weekdays. In this case, we defined weekdays as Monday-Thursday. As seen in Figure 12 below, in our weekday evaluation, we see good performance across facility types, with R^2 values of 0.74 for roads and 0.86 for trails.

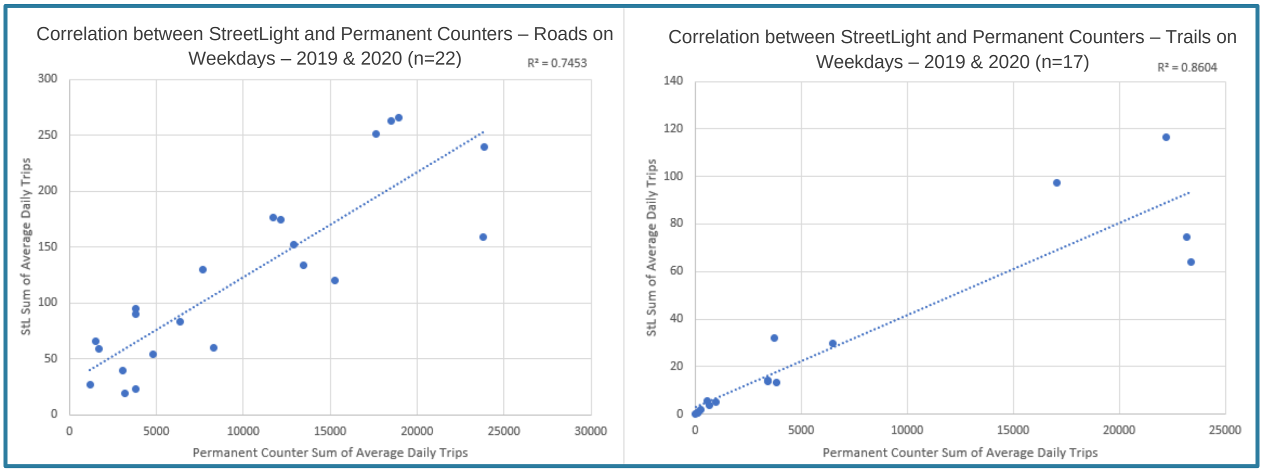

Figure 12: Correlation between average daily bike trips on weekdays. The first chart (left) includes roads only with an R^2 value of 0.74, while the second chart (right) includes trails only with an R^2 value of 0.86. Counter locations are in San Francisco, Philadelphia, and Ottawa.

Overall, given the limited number of permanent counters with data available in 2020, we are very pleased with these results. Customers should note the separate evaluation of roads and trails and emulate this process in their analyses. Specifically, when using reference zones for calibration, or when performing independent validations of the StreetLight Metrics.

Pedestrian Validation

We have taken two primary approaches to validate our pedestrian Metrics. As we did with our bicycle validation, we compare our aggregated trip characteristics to information published by household travel surveys in order to evaluate whether our Metrics are within the general range of expected trip characteristics for each mode. Second, we compare our trip volumes for particular locations (designated pedestrian paths and trails) to published pedestrian counts from permanent pedestrian counters.

We attempt to use only permanent counts that we have deemed reasonable. We understand there is debate around the exact accuracy of permanent count technology (and we found many erroneous points within published pedestrian counts). We do not use temporary or expanded counts for validation, as they have additional sources of error.

Comparing Our Pedestrian Results to Travel Surveys – NHTS and CHTS

We analyzed our aggregated pedestrian trip attributes and compared them to the National Household Travel Survey 2017 (NHTS) and the California Household Travel Survey (CHTS) 2012-2013.

To create a comparison data set, StreetLight Data analyzed 2.1 billion trips across the continental U.S. that occurred between January 2019 and May 2020 for an average of 125.6

million trips per month. NHTS surveyed ~129,000 households across the county and analyzed 81,200 walk trips while CHTS surveyed 42,431 households with fewer than 65,000 walk trips.

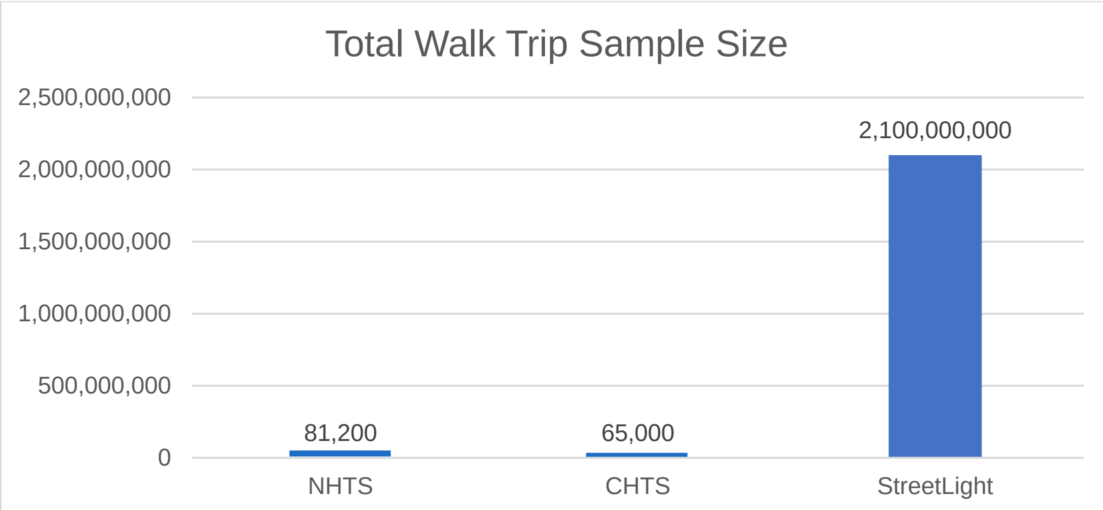

Figure 13: Comparison of NHTS and CHTS total walk trip sample size to StreetLight's walk trip sample size.

We also compared key average characteristics of pedestrian data from our data set to the household surveys. We do not set an "exact match" as the goal. All three data sets are samples, and all three thus will have different strengths and sources of error. Where discrepancies occur, we believe they are explainable by known differences in collection methods or due to known biases in our trip-detection algorithms.

In the Multimode algorithm, current trip boundary detection mechanism requires us to detect a change of mode to terminate a pedestrian trip currently in progress. This algorithm may continue a walk trip a bit further even after the actual trip destination has been reached. For example, in the case of a person walking 10 minutes and then arriving at their workplace, their trip may not end promptly once they arrive at their destination. If they continue to walk around inside their workplace with their device in hand, the trip may take extra minutes to end. This behavior results in trips that are slightly longer in length and duration than survey trips, although the relative locations of the trip origin and destination will be accurate. In the latest iteration of our Metrics, we have made significant enhancements to this "stationary" trip pattern, which has led to improved quality of trip starts and ends, as well as trip duration and speed information. We also expect to see the algorithm capture more short trips than a traditional survey (like the NHTS), as those trips are more likely to be under-reported in traditional surveys. StreetLight currently requires trips to be at least 60 meters long in order to be considered valid.

The following tables show the StreetLight average pedestrian trip attribute Metrics relative to the published NHTS and CHTS numbers.

Table 6: Comparison of average trip length from different sources (miles).

Mode	NHTS	CHTS	StreetLight – National
Walk	0.87	0.30	0.53

As shown in Table 6, we found the StreetLight trip distance averages were within the expected range and consistent with the household surveys. We generally expected StreetLight trips to be shorter in length compared to the NHTS for the following reasons:

1. **Trip-breaking criteria:** StreetLight "breaks" pedestrian trips using a layered approach with multiple conditions. At a high level, if the machine-learning algorithm sees two consecutive points beyond the current trip boundary whose combined probability of walking is < 50%, then we conclude the person may have "changed" modes (which includes coming to a halt since we model "standing still" as a mode in itself). Thus, a walk trip that involves strolling through a park, stopping for lunch along the way, and then continuing to walk afterwards would show up as one long trip in a survey, and two shorter trips in StreetLight.
2. **Recall bias:** Household surveys are based predominantly on recall. Survey respondents are known to be more likely to forget short trips than long ones.[10] This means StreetLight is more likely to pick up shorter trips than the surveys. In particular, the 2017 NHTS moved away from interviewer-assisted recall (which prompted specifically for short and active mode trips), making it more likely that respondents will forget these trips.[11]
3. **GPS assist:** The CHTS used GPS assist to help measure this trip-length bias and found that shorter trips were far more likely to be under-reported than long ones.[12] They then adjusted results to help correct for this bias. For this reason, it's not surprising that the average CHTS length is shorter than the NHTS length.

Figure 14 shows the distribution of pedestrian trips by length. As expected, StreetLight captures many short trips under 0.5 miles, and slightly more long-distance trips than traditional surveys, but overall trends look similar.

[10] Wolf, J., M. Loechl, M. Thompson, and C. Arce (2003). "Trip Rate Analysis in GPS-Enhanced Personal Travel Surveys", in P. Stopher and P. Jones (editors), *Transport Survey Quality and Innovation*, Pergamon Press, pp. 483-498.
[11] McGuckin,N. and A. Fucci. "Summary of Travel Trends: 2017 National Household Travel Survey." FHWA-PL-18-019. July 2018.
[12] Zmud, Johanna. "Identifying the Correlates of Trip Misreporting - Results from the California Statewide Household Travel Survey GPS Study." 10th International Conference on Travel Behavior. August 2003.

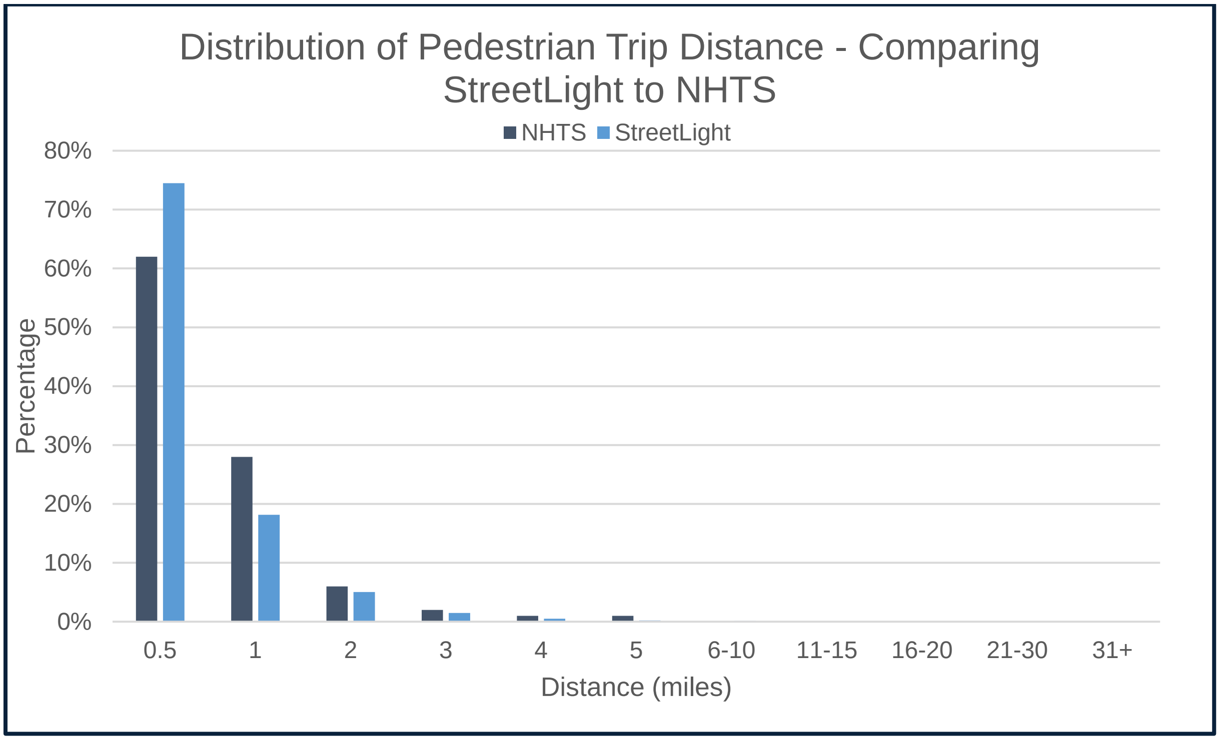

Figure 14: Histogram of pedestrian trip-length distribution for StreetLight and the NHTS.

Next, we compared average trip duration and average trip speed to the survey results. Overall, we observe that the StreetLight average pedestrian trip duration is longer than the reported NHTS and CHTS metrics, as shown in Table 7. However, as discussed previously, we feel this is due to the fact that we continue to log minutes for a trip when the traveler has arrived and is still walking around in the destination building. As a result, the average trip speeds for StreetLight are lower (see Table 8).

Table 7: Comparison of average trip duration (minutes)

Mode	NHTS	CHTS	StreetLight Trips - National
Walk	16.00	10.90	21.79

Table 8: Comparison of average trip speed (mph)

Mode	NHTS - Imputed	CHTS – Imputed	StreetLight Trips - National
Walk (avg. distance / avg. duration)	3.26	1.65	1.77

NHTS documenters reported that recall respondents often report trips in five-minute increments. This could lead to over-estimation of a seven-minute trip as 10 minutes, thus muddying speed and duration statistics, especially for shorter trips. We expect StreetLight walk trip speeds to be slightly slower than average walk speeds because the StreetLight trips include "transition" moments or slowly ending trips that increase trip duration with little movement in distance.

Comparing Our Results to Permanent Pedestrian Counts

To validate data at a more granular level we compared pedestrian volume counts for months in 2019 and 2020 at specific locations to permanent pedestrian counters across cities nationwide.

PEDESTRIAN COUNTER SOURCE REVIEW

For this validation, we looked to analyze permanent pedestrian counters with data available through Spring 2020 in order to be able to validate as many of StreetLight's available months as possible. Since many sources for permanent count data are published annually, this left us with a limited set of locations for comparison. We essentially needed counters that published their data on a daily or monthly basis in order to meet our strict criteria for time coverage. Luckily, we were able to find some high-quality sources in both the U.S. and Canada.

Philadelphia had a number of reliable permanent pedestrian counters covering the January 2019 – May 2020 data period available for analysis. This data is made available by the Delaware Valley Regional Planning Commission (DVRPC)[13]. DVRPC maintains a number of permanent counters across the Philadelphia and the surrounding region. In all, there were 16 permanent counters that collected data in the months we were looking to analyze. Of those 16 counters, only eight provided data consistently across the 17-month period. Three of the 16 permanent counters were located on roads, while the rest were on trails, usually bike and pedestrian paths in parks in Philadelphia and the greater Delaware Valley.

We were also able to source an additional permanent counter in Canada. Specifically, a bike and pedestrian bridge in Ottawa, Ontario known as the Adawe Crossing. This data was made available by the City of Ottawa.[14]

Both data sources reported bike counts hourly over the 17-month period. Thus, we were able to also aggregate daily and monthly totals from the dataset. We performed additional filtering on the monthly counts to ensure that locations and months with wild swings were not included in our comparisons.

In evaluating the permanent counters for comparison, it's first important to consider the range of average daily trip counts recorded across the locations. Due to the limited number of permanent pedestrian counters on roads, we've decided to combine roads and trails in our evaluation.

[13] https://www.dvrpc.org/webmaps/pedbikecounts/
[14] https://open.ottawa.ca/datasets/bicycle-trip-counters

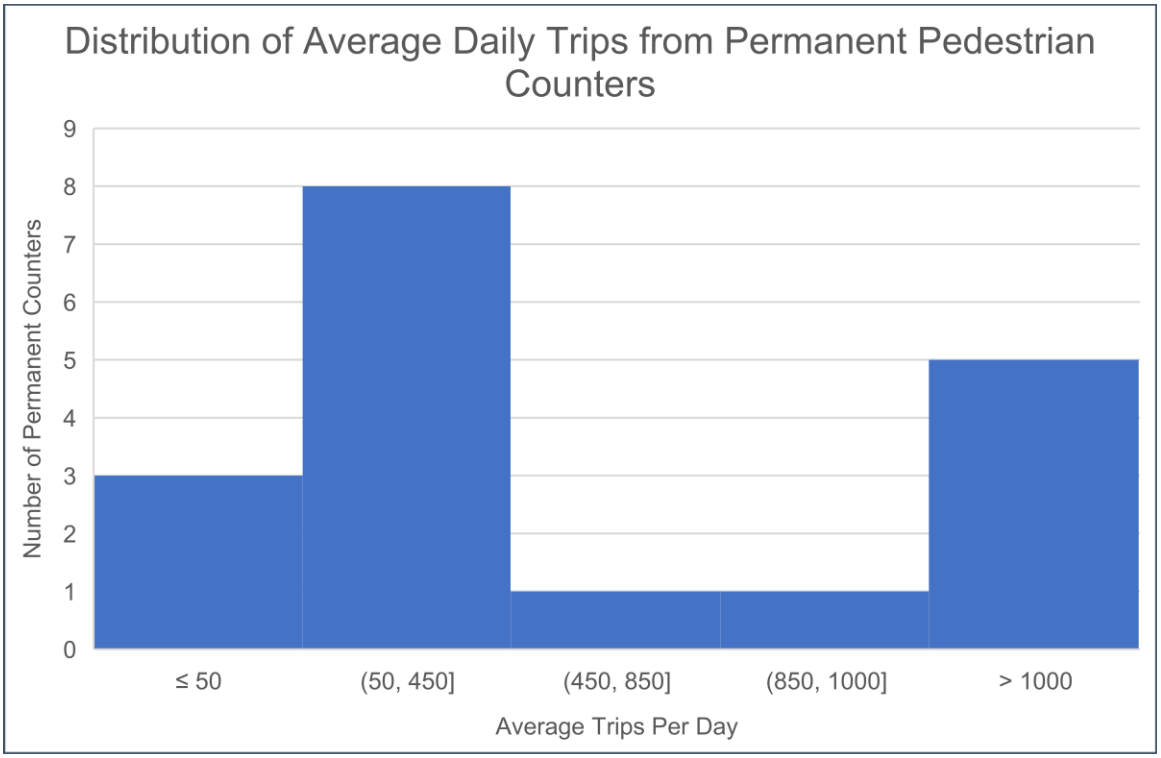

Figure 15: Histogram of selected permanent counters' average trips per day on roads and trails. Most locations record less than 450 trips per day. Counters are located in Philadelphia and Ottawa.

Many permanent counters had relatively low daily trip counts, with daily totals clustered between 50 and 450 trips per day. A few counters, on both roads and high activity trails recorded more than 1000 daily trips on average.

Initial Comparison to StreetLight

After the source data was cleaned, we compared the permanent counts to the StreetLight counts on the matching pedestrian segments obtained from the Streetlight InSight® platform. Our goal was to determine the correlation between the two sets of counts. Since we were working with a limited number of permanent counters, we wanted to include as many months as possible in our comparison, and not necessarily fully exclude counters with missing months, or highly variable months (indicating some counter error). This means that if a counter had months missing, or highly variable outlier months between January 2019 and May 2020, their viable months could still be used for comparison. Therefore, the aggregated comparisons performed for some counters may be the sum of all 17 months, while others may be a sum of a smaller subset of months.

There are a variety of methods to count pedestrians that are commonly classified by technology type and by data collection period. The common types of technology used include passive infrared counters, radio beam, automated video, and manual counts. Almost all these methods can be installed on a temporary or permanent basis and can be distinguished as short duration or continuous counters.

Infrared counters are prone to error and thus require a careful installation process in order to place the equipment at the recommended height level as well as the site-specific location. Counters need to be located in such a way that they only capture pedestrian trips and do not inadvertently capture vehicles or other modes. Additionally, FHWA recommends that a calibration and validation process be conducted on the specific technology to ensure that the count data is within the bounds of acceptable accuracy.[15]

To accurately count only pedestrians on facilities such as Class I multi-use paths, additional counters must be installed at the same location that can count bikes so that bike volume can be subtracted from the pedestrian counts. This typically involves installing loop detectors or pneumatic hose counters to capture bike traffic in one direction on the road along with an infrared counter to capture pedestrian trips. The total volume from the infrared counter is tabulated, and pedestrian trips are subtracted from the bicycle counts.[16]

StreetLight's data for counter-matching has an additional source of potential error: trip "locking," or assigning a pedestrian trip to a particular route or path, based on waypoints measured during the trip. This is a bigger source of error for pedestrians than for cars or bicycles, as pedestrians often operate without following rules. Pedestrian movement is more fluid than vehicular or bicycle travel patterns — pedestrians can cross streets mid-block, cut-through buildings, and move across non-designated pathways, ultimately creating more challenges in the trip-locking process. Thus, we expect our validation results for counter-matching to be weaker than validation to aggregate trip metrics, or to an Origin-Destination study that looks just at trip starts and stops.

In order to get a sense of our general performance across counter locations, it's important to consider the rates of penetration recorded at each counter. Penetration rates are calculated as the sum of the monthly StreetLight sample divided by the sum of monthly counter sample. A relatively consistent penetration rate is key to consistent results in a validation. Penetration rates are also key to helping us understand which counters and months are outliers, indicating some other source of error may be at play.

[15] FHWA Traffic Monitoring Guide 2016 https://www.fhwa.dot.gov/policyinformation/tmguide/
[16] National Cooperative Highway Research Program (NCHRP) Report 797: Guidebook on Pedestrian and Bicycle Volume Data Collection http://www.trb.org/Main/Blurbs/171973.aspx

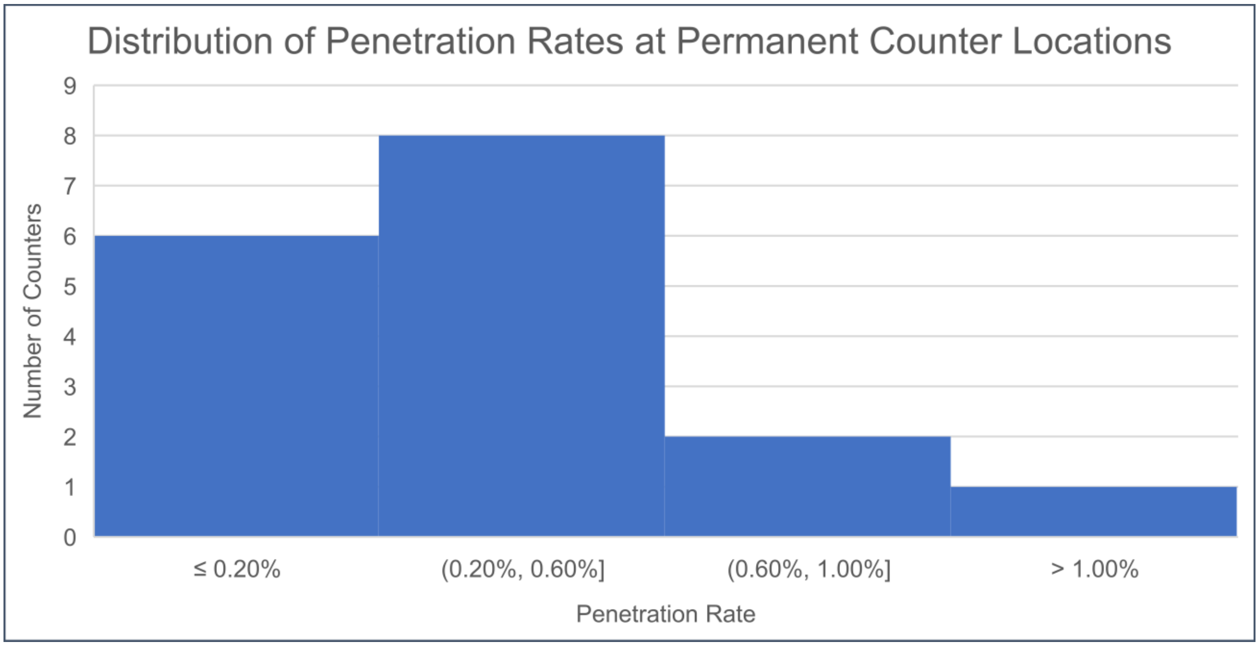

Figure 16: Distribution of penetration rates for selected locations on roads and trails. Most locations see penetration rates between 0.20 and 0.60%. Counters are located in Philadelphia and Ottawa.

Generally, average penetration rates for pedestrian trips are clustered around 0.50%. Given the fact that most counters are on trails with relatively low sample sizes, penetration rates below 1% are expected. If we were to obtain additional counts on roads, we would expect to see higher penetration rates (as we do with bike trips in urban areas).

TIME TRENDS

In validating our results, we first decided to evaluate hourly distribution across average weekends and weekdays. For the day type and day part comparison, we first looked at the sum of trips across hours for all locations to see if the distribution was consistent.

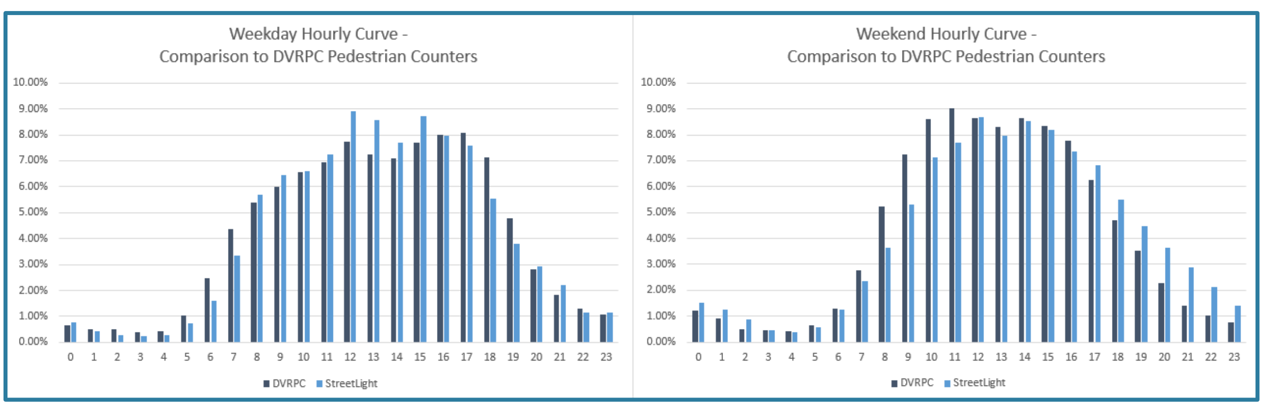

Figure 17: Average hourly curves on weekdays and weekends, comparing DVRPC's permanent counters to StreetLight Metrics at the same locations between Jan 2019 and May 2020.

With zones primarily on trails, the DVRPC locations show a slight morning and afternoon peak, though a less normal curve than the weekend curve. StreetLight appears to show a slightly higher mid-day peak during weekdays, though the shape of the curve is similar overall. Our ability to mirror these unique hourly trends confirms our confidence in comparing weekdays to weekends, as well as comparing day parts to one another. In general, but on weekends in particular, StreetLight under-samples compared to during the mornings, compared to early afternoons. This requires further investigation and potential algorithm adjustments, but may be also accounted for in factoring. Overall, the individual count locations showed similar trends compared to the aggregate.

In addition to evaluating hourly trends, we also wanted to confirm that our normalization process enabled users to capture monthly trends. Our normalization process, which uses monthly scaling values derived from comparisons to permanent counters, should ideally capture the seasonal or monthly curves recorded by the permanent counters. Since we apply a single scaling value to each month (for more information see section on normalization in our earlier methodology description), we are not expecting to produce estimated counts that perfectly mirror the monthly permanent counts. However, we are looking to capture the variation across time so that months can be directly compared to one another using the normalized StreetLight Index value. So as not to bias the results, none of the locations listed below in Figure 18 were utilized in the normalization process.

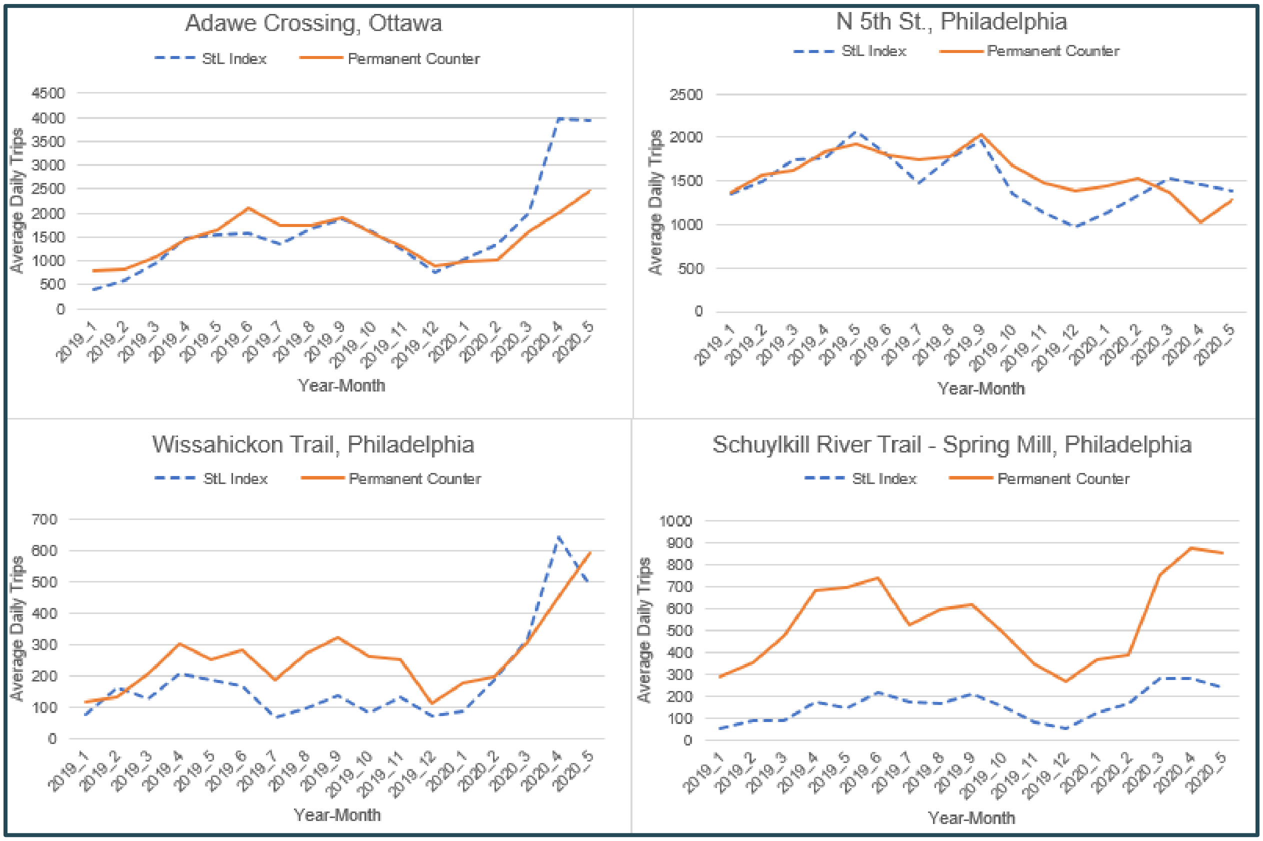

Figure 18: Average monthly curves, comparing permanent counter locations in Philadelphia and Ottawa to the StreetLight Index between Jan 2019 and May 2020. The year and month combination is represented on the x-axis, while the average daily trips are represented on the y-axis.

The charts above indicate that the StreetLight Index is able to capture seasonal and monthly curves reported by the permanent counters. Note that the N 5th St. example in Philadelphia is one of the few pedestrian counters on a road as opposed to a trail, and we see strong results there. We're pleased that the Index is able to capture seasonal peaks, as well as changes during 2020 months impacted by COVID-19 (March through May, 2020).

COMBINED CORRELATION RESULTS

To evaluate correlation between counter locations, we combined results from Philadelphia and City of Ottawa. First, we evaluated the correlation between StreetLight counts and the sum of average daily trips, recorded by month, for each permanent counter. For counter and month combinations with penetration rates outside of our expectations, we removed those months as outliers, and summed the remaining non-outlier months for the given counter. This included removing months with penetration rates above 2% and below 0.10%. Many trails had very low monthly counts, thus we want to ensure a sufficient monthly sample. This allowed us to handle cases where error may have been introduced into a single month, without impacting all remaining months.

We evaluated results across trails including all available 2019 and 2020 months. Results were strong with an R^2 value 0.71. This means that StreetLight's sample alone "explains" 71% of the variation in the permanent counters on the selected trails across 2019 and 2020. No factoring or adjustment has been done to these results. The FHWA recommends that even permanent counters be factored based on location type, sensor type, season/weather, and more to compensate for inaccuracies.[17]

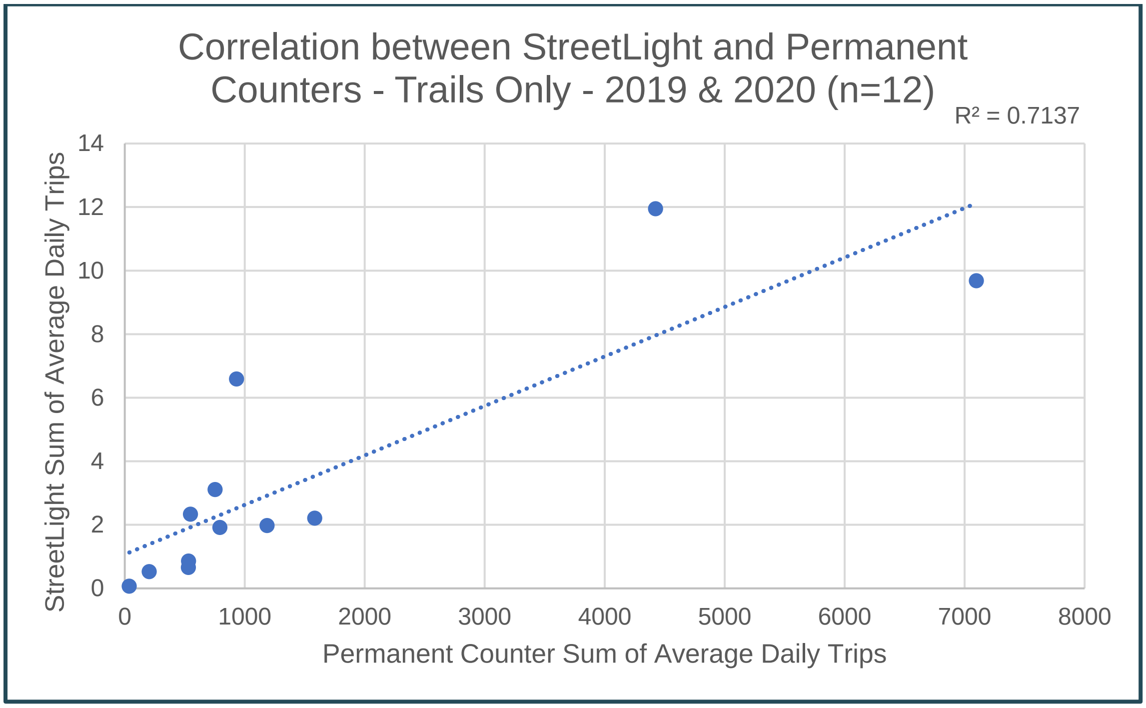

Figure 19: Correlation between average daily pedestrian trips reported by permanent counters and average daily bike trips reported by StreetLight on trails.

Overall, given the limited number of permanent counters with data available in 2020, we are pleased with these results. While we didn't have a sufficient volume of permanent counters on urban roads or sidewalks to properly evaluate those trends, it is reasonable to assume that, like our bike Metrics, roads and trails will see different rates of penetration and thus should be evaluated separately. We are seeking additional permanent pedestrian counts on roads in future iterations of this validation document. Going forward, customers should note the separate evaluation of roads and trails, and emulate this when using reference zones for calibration, or when performing independent validations of the StreetLight Metrics.

About StreetLight Data

[StreetLight Data, Inc.](#) pioneered the use of Big Data analytics to help transportation professionals solve their biggest problems. Applying proprietary machine-learning algorithms to over four trillion spatial data points over time, StreetLight measures multimodal travel patterns

[17] https://www.fhwa.dot.gov/policyinformation/travel_monitoring/pubs/pedbikedata.cfm#sect7_3

and makes them available on-demand via the world’s first SaaS platform for mobility, StreetLight InSight®. From identifying sources of congestion to optimizing new infrastructure to planning for autonomous vehicles, StreetLight powers more than 6,000 global projects every month.

STREETLIGHT DATA