STREETLIGHT InSight

StreetLight AADT 2020 Methodology and Validation White Paper

Canada

Version 1.0

August 2021

Table of Contents

Introduction: Why We Developed an AADT Metric for Canada

Annual Average Daily Traffic (AADT) volumes are a core metric in transportation practice. Due to the lack of up-to-date AADT counts in many Canadian provinces, many planners don't have access to up-to-date traffic volumes on their roads, which are needed for transportation planning. Understanding traffic volume can help drive infrastructure and planning decisions and influence policy. This lack of access to traffic volumes has increased the need for quick, easy, and cost-effective AADT measures, and for related metrics such as Vehicle Kilometers Traveled or "VKT." Traditional methods to estimate AADT, such as temporary count expansion, have relatively large errors and are cumbersome and expensive to deploy. In order to help agencies and practitioners overcome these challenges, StreetLight Data has developed a comprehensive AADT Metric for more than a million kilometers of urban and rural roadway in Canada. It outperforms industry-standard accuracy targets and can be used by industry practitioners for a wide variety of applications.

What's New in AADT 2020 for Canada

StreetLight's AADT 2020 Metric uses a modelling framework that is similar to the Canada AADT 2019 V3. Some highlights and improvements of the new AADT Metric are:

- StreetLight has access to training data for the years 2018 through 2020 from a total of 2,510 permanent counters (in 1,974 unique locations) across nine provinces, which is ~200 more counters than AADT 2019 V3. It should be noted that not every counter has available data for all three years.
- For the year 2020, training data is available from 1,100 permanent counters (in 747 unique locations) located in the four provinces of British Columbia, Alberta, Manitoba, and Nova Scotia. Results are optimized for these provinces, and we plan to enhance the model when data from the remaining provinces becomes available.
- The inclusion of all historic training data (from 2018 and 2019) in the model allows us to produce better AADT 2020 estimates, even for provinces where training data is not available in 2020.
- A new type of "hybrid" machine-learning model was implemented in each province, improving estimation on low- and medium-volume roads across the provinces.
- Features (or variables) were tuned to improve the model performance.

Defining Target Accuracy

Our first task was to define what was "accurate enough" for our AADT estimate. We had to balance the need for accuracy with the need to build an algorithm that could calculate an AADT very quickly while being compatible with the StreetLight InSight® transportation analytics platform. This meant that the algorithm had to be computationally elegant and scalable.

To evaluate accuracy, we grouped permanent counter stations by AADT, and reported on results within those groupings (labeled as road-size bins). On some occasions, we used 10 road-size bins for more detailed results, while in others we used three road-size bins to more generally describe high-, medium-, and low-volume roads. Errors measured as a percentage of AADT naturally increase with smaller roads, and groupings by road AADT allow for more visibility into where the errors of the model lie. For summary metrics to describe the general error across an AADT grouping, we report on mean absolute percentage error (MAPE; see equation below), and normalized root mean square error (NRMSE, see equation below). These metrics are standard measures of error for predictive machine-learning models and are helpful to give an idea of what an expected error is for most roads. We detail both metrics because they are both commonly used in literature and highlight different aspects of model performance. MAPE better describes errors on small roads and treats errors of all sizes equally, and in contrast, NRMSE penalizes large errors more, making it more sensitive to the accuracy of AADT estimation on high-volume roads.

The equation for mean absolute percentage error (MAPE) is shown below:

$$MAPE \ = 100 * \left(\frac{1}{n}\right) * \sum_{i=1}^{n} \left| \frac{AADT_{Estimate(i)} - \ AADT_{Permanent\ counter(i)}}{AADT_{Permanent\ counter(i)}} \right|$$

The equation for normalized root mean square error (NRMSE) is shown below:

$$NRMSE \ = 100 * \frac{\sqrt{\left(\frac{1}{n}\right) * \ \sum_{i=1}^{n} \left(\ AADT_{Permanent\ counter(i)} \ - AADT_{Estimate(i)} \right)^2}}{\left(\frac{1}{n} * \sum_{i=1}^{n} (AADT_{Permanent\ counter(i)} \ \right)}$$

While the summary of errors is useful for getting a sense of model performance, we find that describing the percentile spread of error across the test set is additionally illustrative of where error lies. The 68th percentile absolute error is reported, as it represents one standard deviation from the mean within a standard bell curve, and thus is a useful descriptor of the typical error across the road segments. As MAPE and NRMSE may be sensitive to outliers, the 68th percentile absolute error can provide more visibility into expected “typical” error. The 95th percentile absolute error is provided to measure the spread of errors across a broader array of sites, representing the upper limit of the expected absolute error. Finally, we present the median percentage error (note that this is not absolute error) to give an indication of bias in our model. The median is chosen instead of the mean because it is less affected by the often right-skewed distribution of percentage error, which is caused by the restriction that AADT estimates cannot be negative. The median is also less sensitive to extreme outliers, which may occur due to errors independent of the model, such as map matching errors in the trip routing algorithms or errors in reference permanent counts. Values close to zero suggest that our model has low bias, while positive values indicate overestimation, and negative values indicate underestimation.

Calculating the Accuracy of StreetLight's AADT 2020 Metric

StreetLight's AADT 2020 model incorporates historic training data from 2018 and 2019 (available in nine provinces) along with 2020 training data (available in four provinces) to produce better estimates across the provinces. In total, we trained our model on data from 2,510 permanent counters (in 1,974 unique locations). The leave-one-out cross-validation (CV) technique is used to evaluate the model performance. That means we train on 1,973 locations and test the model performance on the remaining location. This process is repeated 1,974 times so that all counters in each location have a chance to be in the test set. We use the leave-one-location-out technique instead of leave-one-counter-out for CV to prevent model leakage (i.e., we do not train on one counter from a location and test on the other). In this section and subsequent ones, we present the model performance on the test set with 2020 data. That includes data from the four provinces of Alberta, British Columbia, Manitoba, and Nova Scotia only.

As shown in Figure 1, the actual and estimated AADT values through the cross validation are highly correlated with an extremely high R^2, and the fitted line is very close to a 45-degree line, which indicates the strong performance and low bias of our model.

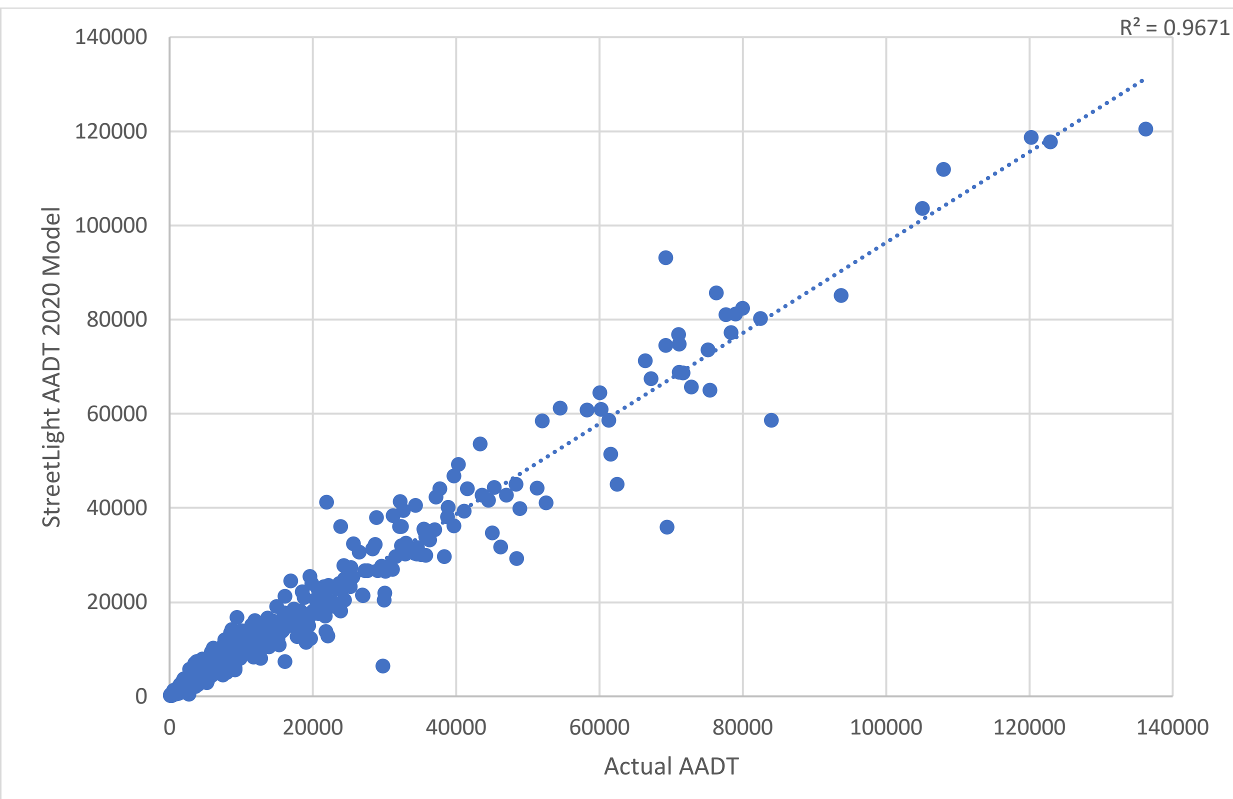

Figure 1: StreetLight AADT 2020 for test data compared to permanent counter AADT. Results are shown only for Alberta, British Columbia, Manitoba, and Nova Scotia.

Table 1 shows the estimation accuracy using the leave-one-out CV method for the eight categories of road size (A-H) in order to gain further insight into the model's performance. We can see that the model performs well, especially on higher-volume roads, and has minimal bias (as shown in the 50th percentile column) across the road size categories.

Road Size	Count	Median Bias	68th Abs Percentile	95th Abs Percentile	MAPE (%)	NRMSE (%)
A: 0 - 499	69	11.58	30.60	62.32	25.59	32.82
B: 500 - 1,999	288	1.66	21.29	45.40	17.38	25.21
C: 2,000 - 4,999	285	0.28	17.16	44.83	16.26	24.81
D: 5,000 - 9,999	182	-1.14	15.38	44.28	13.86	20.81
E: 10,000 - 19,999	150	-1.67	12.29	31.81	11.44	16.41
F: 20,000 - 34,999	62	-2.71	15.62	42.07	15.33	22.03
G: 35,000 - 54,999	33	-4.34	16.03	26.64	12.02	16.31
H: 55,000 +	31	-1.46	8.53	32.38	9.50	13.48

Table 1: Leave-one-out cross validation results for StreetLight's AADT 2020 Metric. Results are shown only for Alberta, British Columbia, Manitoba, and Nova Scotia.

In Table 2, we present the overall estimation accuracy from the CV results for each of the four provinces used in the training data set. Of these provinces, Alberta and British Columbia contribute the most training data and are in the top four biggest provinces in Canada by population. Hence, it is important for us to improve the estimation accuracy across different road sizes for them.

Province	Count	Median Bias	68th Abs Percentile	95th Abs Percentile	MAPE (%)	NRMSE (%)
Alberta	708	0.25	16.70	39.53	14.34	24.61
British Columbia	345	-0.54	20.12	58.38	18.54	34.36
Manitoba	14	-5.30	12.33	17.77	9.27	13.30
Nova Scotia	33	0.10	23.81	43.36	18.45	28.11

Table 2: Leave-one-out cross validation results for StreetLight's AADT 2020 Metric for different provinces used in machine-learning model training.

Tables 3a and 3b present the estimation accuracy for the four road size categories (very low, low, medium, and high) in the provinces of Alberta and British Columbia for AADT 2020.

Road Size	Count	Median Bias	68th Abs Percentile	95th Abs Percentile	MAPE (%)	NRMSE (%)
0 – 499 (very low)	69	11.58	30.60	62.32	25.59	32.82
500 - 4,999 (low)	427	1.28	17.83	36.98	14.49	19.82
5,000 - 54,999 (medium)	193	-2.09	12.15	30.38	10.78	18.83
55,000 + (high)	19	-2.20	7.62	14.00	6.19	7.63

Table 3a: Leave-one-out cross validation results for StreetLight's AADT 2020 Metric for Alberta.

Road Size	Count	Median Bias	68th Abs Percentile	95th Abs Percentile	MAPE (%)	NRMSE (%)
0 – 499 (very low)	N/A	N/A	N/A	N/A	N/A	N/A
500 - 4,999 (low)	126	0.72	27.05	73.59	24.28	32.99
5,000 - 54,999 (medium)	207	-1.19	17.82	44.32	15.27	27.17
55,000 + (high)	12	1.91	19.85	40.76	14.73	20.66

Table 3b: Leave-one-out cross validation results for StreetLight's AADT 2020 Metric for British Columbia.

Overall, we are pleased with the results in the provinces used for training. Note that we do not have the training data for 2020 for the two big provinces of Ontario and Quebec, and we expect that the model would have produced slightly higher errors for those two provinces than it does for Alberta and British Columbia. However, the inclusion of historic data for Ontario and Quebec in the model helps improve the accuracy of the AADT 2020 estimates, more than it would have otherwise. We refer readers to our Canada AADT 2019 V3 whitepaper to see the expected results in Ontario and Quebec for 2019 AADT, and we expect the 2020 AADT results for Ontario and Quebec to have just slightly higher errors than those.

Comparing Cross-Validation Results between AADT 2019 and 2020

In this section, we compare the cross-validation results between Streetlight AADT 2019 V3 and AADT 2020 Metrics for the two large provinces of British Columbia and Alberta in Tables 4a and 4b. Reference permanent counts for 2020 were only available for the four provinces of British Columbia, Alberta, Manitoba, and Nova Scotia. Analysis of model error at permanent counter locations within the two large provinces of British Columba and Alberta allows for direct comparison of differences between the two models.

It can be seen that except for the very low volume category (0-499) in Alberta, where the 2020 model has higher median bias than the 2019 V3, the two models have comparable performance. This is due to their similar modelling frameworks which optimize the model performance in each province. This indicates that despite dramatic changes in travel patterns due to COVID-19, our estimations remain robust in 2020.

Road Size	Count	Median Bias (%)		MAPE (%)	
		StreetLight 2020	StreetLight 2019 V3	StreetLight 2020	StreetLight 2019 V3
0 – 499 (very low)	69	11.58	6.92	25.59	23.97
500 - 4,999 (low)	427	1.28	-0.38	14.49	15.86
5,000 - 54,999 (medium)	193	-2.09	1.70	10.78	13.85
55,000 + (high)	19	-2.20	-2.08	6.19	7.68

Table 4a: Comparing the cross-validation results between the AADT 2020 and 2019 V3 models for Alberta

Road Size	Count	Median Bias (%)		MAPE (%)	
		StreetLight 2020	StreetLight 2019 V3	StreetLight 2020	StreetLight 2019 V3
0 – 499 (very low)	N/A	N/A	N/A	N/A	N/A
500 - 4,999 (low)	126	0.72	1.45	24.28	26.56
5,000 - 54,999 (medium)	207	-1.19	-2.61	15.27	16.41
55,000 + (high)	12	1.91	1.86	14.73	20.14

Table 4b: Comparing the cross-validation results between the AADT 2020 and 2019 V3 models for British Columbia

Data Sources Used in Our AADT 2020 Metric

Our AADT 2020 considers the following data sources when developing features to create our best prediction of Annual Average Daily Traffic:

Input 1: Location-Based Services Trips Data

Location-based services (LBS) data is created by smartphone applications providing a service that depends upon on a device's geographic location in the physical world — for example, shopping apps, weather apps, or dating apps. We use algorithmic processing techniques to link these data points into trips. We sampled trips throughout 2020 in order to create the best model possible for the 2020 calendar year.

Input 2: Navigation GPS Trips – Commercial

The navigation-GPS data we use is created by connected commercial vehicles. Our data set is tagged by vehicle type: heavy-duty commercial vehicle or medium-duty commercial vehicle. Since roads vary heavily in the share of commercial trucks (and in the share of medium vs. heavy-duty trucks), having a combination of data sources from commercial vehicles is critical.

Input 3: Canadian Census Data

We normalized our LBS trips using data from Statistics Canada. Normalizing is an important step to adjust a sample that is not perfectly distributed. In short, if 10 devices in our sample "live" in a dissemination area (DA) with 100 people, each of those devices is scaled up by a factor of 10. If ten devices "live" in a DA with 50 people, each is scaled by a factor of five. This adjusts for variation in geographic distribution, which is correlated with demographic factors like income. We also looked at the population density near the road in question (which is a proxy for identifying a road as rural or urban), as well as income data pertinent to the surrounding area.

Input 4: OpenStreetMap Data

We included features commonly extractable from OpenStreetMap (OSM) such as road geography, speed limits, number of lanes, availability of parking, road classification, and other factors. We know that all OSM features are not always available for every road. Our algorithm is factored to adjust to a different set of coefficients if no OSM feature data is available. We also use the OSM to “lock” a trip to a route by connecting pings along the most viable network path a vehicle can take.

Input 5: Weather Data

We included data on precipitation and temperature to account for areas that have extreme precipitation events (like snow storms) on a regular basis and might experience different travel patterns as a result.

Input 6: Training and Testing AADT Using Permanent Loop Counters

We researched extensively to find well-cleaned permanent counter data. We wanted our data to be spread across Canada, between small and large roads, urban and rural. The biggest challenge was finding permanent counter data for small rural roads. Below we present the map showing the locations of the 2,510 permanent counters in Canada for which we have training data. The number of permanent counters used in each province is shown in Table 5.

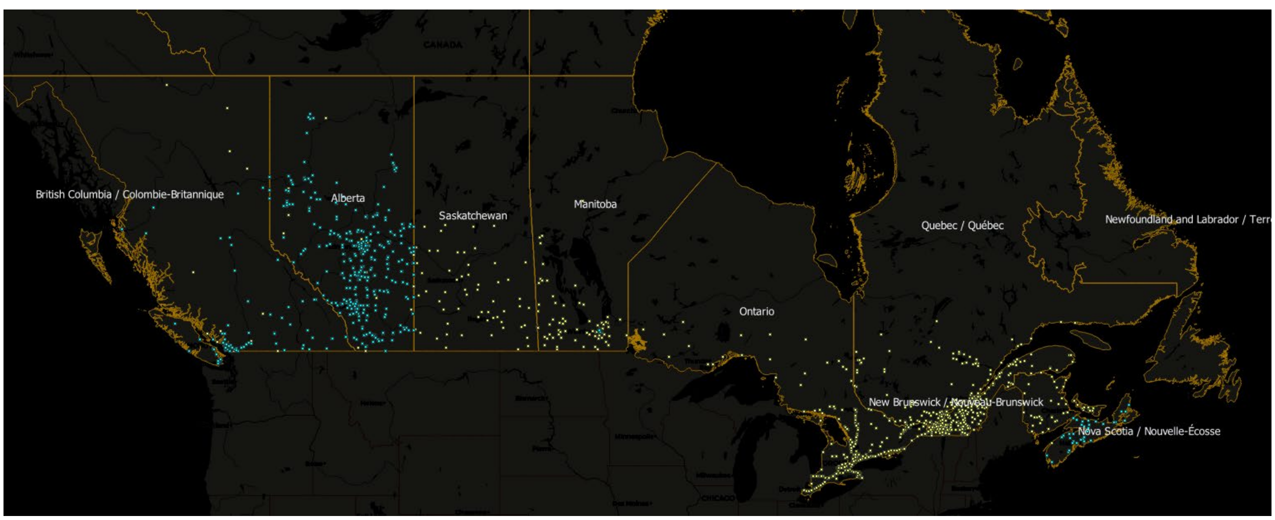

Figure 2: Map of all permanent counters used for training and testing the AADT 2020 Metric. (The turquoise counters are those with 2020 training data and the yellow ones are those with historic data from 2018 and 2019.)

Province	# Permanent Counters
Alberta	827 (362)

British Columbia	592 (345)
Manitoba	171 (7)
New Brunswick	36 (0)
Nova Scotia	34 (33)
Ontario	404 (0)
Prince Edward Island	11 (0)
Quebec	376 (0)
Saskatchewan	59 (0)

Table 5: Number of permanent counters across nine provinces that are used for training and testing the AADT 2020 Metric. In brackets are the number of permanent counters with training data in 2020.

Road Size	# Permanent Counters
A: <= 499	147 (59)
B: 500 - 1,999	576 (219)
C: 2,000 - 4,999	575 (263)
D: 5,000 - 9,999	444 (182)
E: 10,000 - 19,999	403 (150)
F: 20,000 - 34,999	217 (62)
G: 35,000 - 54,999	106 (33)
H: 55,000 - 84,999	73 (31)
I: 85,000 - 124,999	66 (0)
J: 125,000+	88 (0)

Table 6: AADT range for all permanent counters used for training and testing the AADT 2020 Metric. In brackets are the number of permanent counters with training data in 2020.

Selecting and Testing the Algorithm

We weighed dozens of different algorithmic approaches when developing our AADT 2020 Metric. In this section, we will provide an overview of our major decisions.

First, we checked to see how our normalized LBS trips, which comprise our strongest and largest data set, were correlated with AADT. The results are shown below. As you can see, the correlation is fairly strong. As such, the remainder of the task was to use machine learning to reduce error and improve correlation (compared to scaling to counts from LBS trips alone).

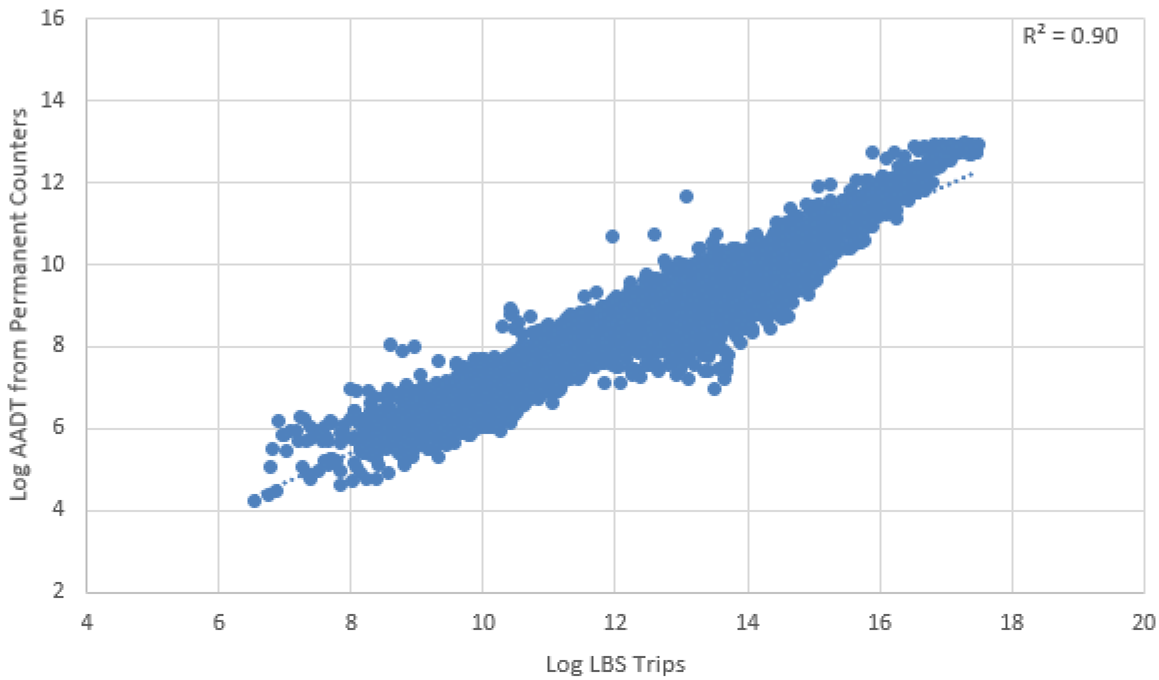

Figure 3: Correlation of population-normalized LBS to permanent loop counter data – the single best predictor value. The rest of the machine-learning work aimed at improving these results.

The choice came down to three options for machine-learning techniques: ordinary least squares (OLS), random forest, and gradient boosting.

First, we tried OLS, a multivariate equation framework with machine learning. The benefit of a multivariate regression technique is that it is easier to explain, as it is more or less building a classic y = mx + b style equation. The disadvantage, as we've found, is that the results were not as accurate as we wanted, and the model was prone to throwing outliers.

We also tried a random-forest model, which we've relied on in our prior years' versions of our AADT models. The benefit of a random-forest model for AADT estimation is that it did a far better job of handling unusual roads (such as small ones or those with extremely high commercial traffic near ports and warehouses). The disadvantage is that the algorithm is less intuitive than a regression technique when explaining to non-data scientists. In the end, we decided that the accuracy and algorithmic robustness for unusual roads and outliers were more important.

Finally, we explored gradient-boosting and extreme gradient-boosting models, which are also tree-based models, like random forest. We found that in most cases, these options have a narrower spread of errors and lower NRMSE compared to the random-forest algorithm. The gradient-boosting algorithm adds additional complexity to the random-forest algorithm by fitting errors as the model is built, which can further boost model performance. Extreme gradient boosting has faster model run-time than traditional gradient-boosting algorithms, and because the model error was similar, we selected the extreme gradient-boosting model as the final model for AADT 2020 estimation. For more information on extreme gradient boosting, see our summary in Box 1 below.

Box 1: The Extreme Gradient-Boosting Model

An Extreme Gradient Boosting model is a specific implementation of the gradient boosting method. This approach is similar to a decision tree, but it uses several decision trees. For example, let's predict whether a patient entering an emergency room is at high risk. A decision tree may look like this: If age is over 50, blood pressure is over 150, and temperature is above 98 degrees Fahrenheit, then the patient is high-risk. That's a decision tree. It is very interpretative but does not have much predictive power alone. Gradient boosting uses a lot of decision trees (say, an ensemble), where each tree is a little bit different from the others. When a new patient arrives, we take the majority vote of the decision-tree ensemble to get a final result. Gradient-boosting models build each tree one at a time, while alternatives like random-forest models build each tree independently.

The different trees use random samples of observations and subsets of features to train. For example, instead of considering age, blood pressure, and temperature, we may train one tree with age and blood pressure, another with blood pressure and temperature, another with age and temperature, and so on if we had more features. The key is that the trees become a bit different (less correlated), so when the results are combined, we get a "diverse" answer. The idea behind this model is that a bunch of poor decision-makers put together in a room to form a committee will start making better decisions. If each decision-maker comes with a different perspective, that creates better results.

In AADT 2020, we adopted a hybrid-model approach where extreme gradient-boosting serves as the main algorithm and produces the bulk of the results. In a small number of cases (e.g., unusual roads), the random-forest algorithm is utilized. Together, our new hybrid approach boosts the overall estimation accuracy for all road size categories, especially low-volume roads, as well as improves accuracy across the provinces.

The next step was to decide which features (input variables) to include in our models. We tested hundreds of combinations. We wanted results with the least amount of error, but we also wanted our algorithm to be scalable to anywhere in Canada, and to be computationally efficient. In the end, we built four models that relied on a combination of 46 features captured from different data sources. We strive to avoid over-fitting by throwing far too many features into a machine-learning model. This may make initial results look very good, but it also prevents the approach from scaling well outside of the research setting.

Figure 4 below illustrates a high-level flow diagram of how Big Data, specifically Location-Based Services (LBS) and navigation GPS data, and machine-learning models can be used to estimate AADT. First, the machine-learning model is trained to learn the relationship between the AADT derived from permanent counts and Big Data, along with contextual features influencing the AADT. Next, the hyperparameters of the model are tuned through cross-validation to enhance the model performance and avoid overfitting. Finally, the model is applied to new stations (that the model was not trained on) with the input Big Data and contextual features and produces the estimated AADT.

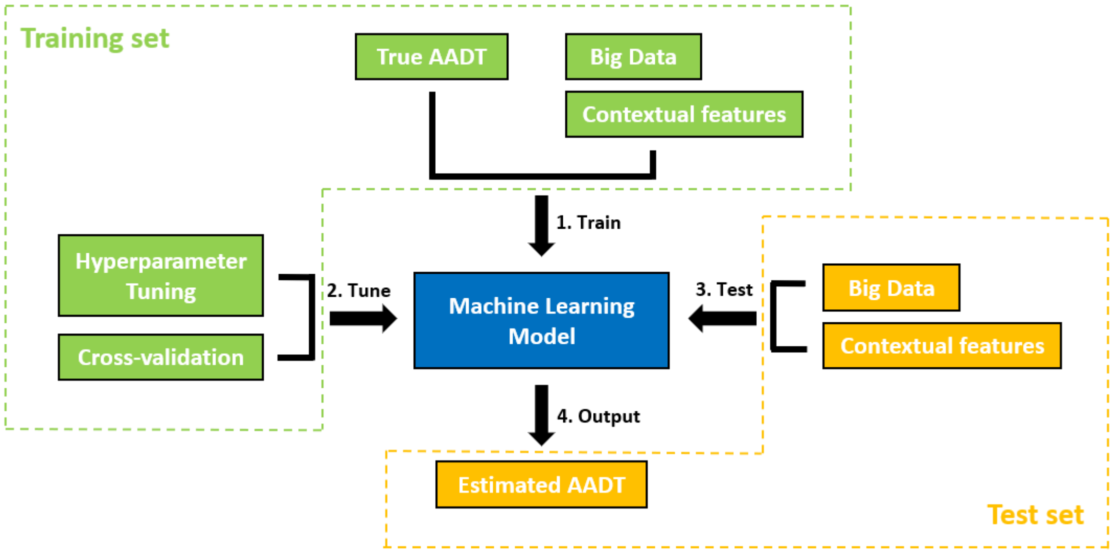

Figure 4: Flow diagram of AADT estimation using Big Data and machine learning.

Our AADT 2020 Metric in the StreetLight InSight® platform also includes a 90% confidence range for each AADT estimate to help locate the true AADT value. Our confidence range (also known as prediction interval) is an estimate of the interval within which the true AADT is expected to lie 90% of the time. To estimate the confidence range, the percent error and log of the predicted AADT values from the cross-validated data set were fit to a quantile regression. For each permanent counter in the cross-validated dataset, the predicted AADT value was mapped to the quantile regression, to determine the upper and lower confidence range limit for a 90% percentile confidence range. Figure 5 is a visual representation of the confidence ranges as they apply to high-, medium- and low-volume roads.

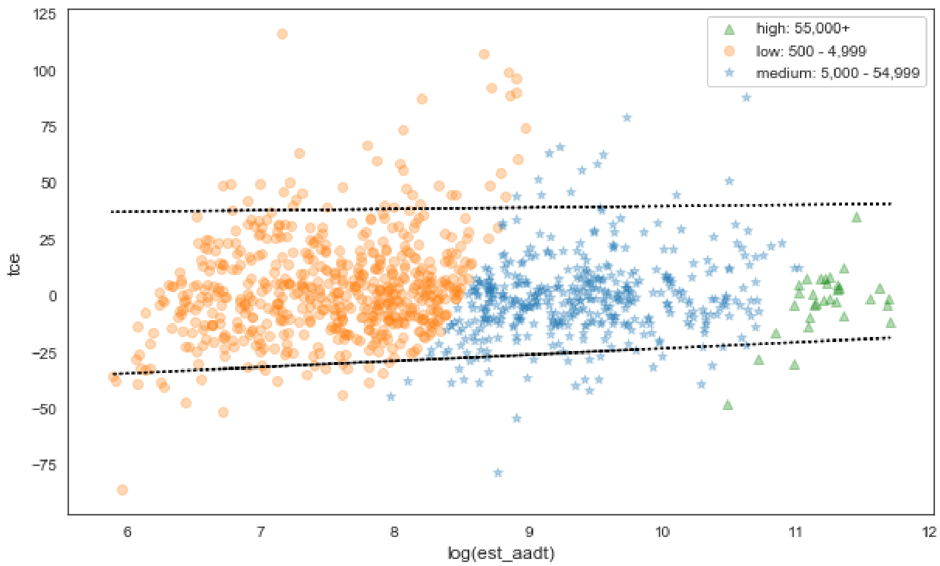

Figure 5: Quantile regression plot of the AADT 2020 model error across three road size bins.

Running AADT 2020 in StreetLight InSight®

If you have a StreetLight InSight account with AADT enabled, choose to create a new analysis. Click the "Create Analysis" button under the AADT analysis option. Name the analysis, select the zone sets covering the roads of interest, and choose the AADT year "2020." Then click "Confirm Analysis" to begin processing.

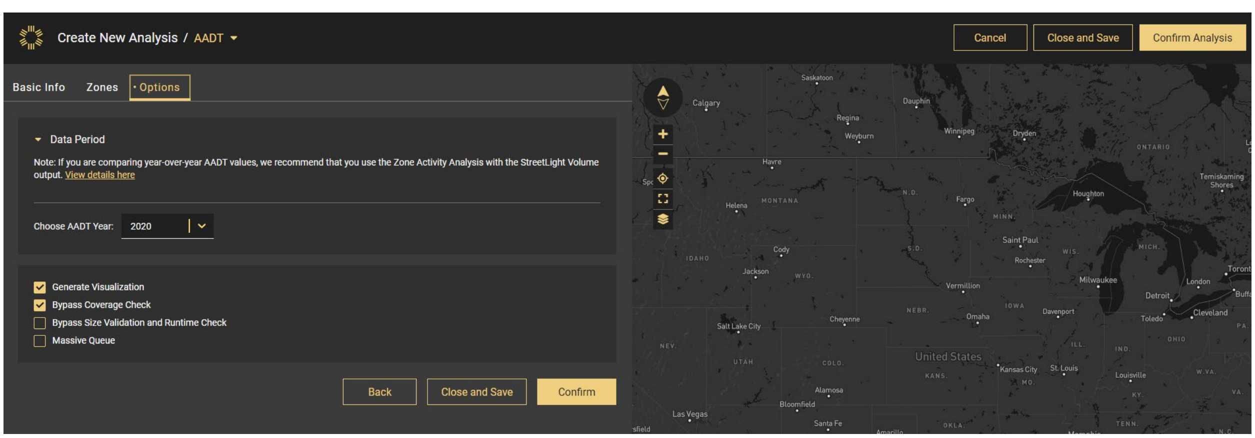

Figure 6: AADT 2020 as represented in the StreetLight InSight® software platform.

Our Next Steps

In future StreetLight InSight® releases, as new data for other provinces become available, StreetLight will update the Canada AADT 2020 model. We'll also work to improve processing speed so that AADT analyses complete more quickly within the platform. In addition, we will continue conducting validation and improvement studies on our AADT metrics.

About StreetLight Data

StreetLight Data, Inc. ("StreetLight") 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 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