University of St.Gallen

Commercial Vehicle Traffic Detection from Satellite Imagery with Deep Learning

Moritz Blattner, University of St. Gallen
Michael Mommert, University of St. Gallen
Damian Borth, University of St. Gallen

Tackling Climate Change with Machine Learning workshop at ICML 2021
July 23, 2021

"From insight to impact"

Introduction

Annual shipping volumes in 2019 passed the 100 billion parcels and is expected to double within the next five years [1]. Road freight is and will remain the dominant mode of surface freight transport and already today **accounts with 2.4Gt for more than 7% of the annual CO_2 emissions** [2,3]. Reliable and granular traffic census is essential for

- Road maintenance and congestion planning.
- Economic analyses and measuring road freight activities.
- Policies for mitigating climate change.

While **carbon-neutral solutions** for heavy modes of transport **have not yet matured**, this can be expected to increase [4].

Most traffic data collection means require physical measurement stations. However, dense **ground-based monitoring systems are costly** and thus often **unavailable in developing countries**, where much of the growth in commercial vehicle traffic is occurring [5]. Around half of all countries worldwide do not maintain traffic census methods and these blind spots **prevent a consistent and widespread analysis of road freight activity**.

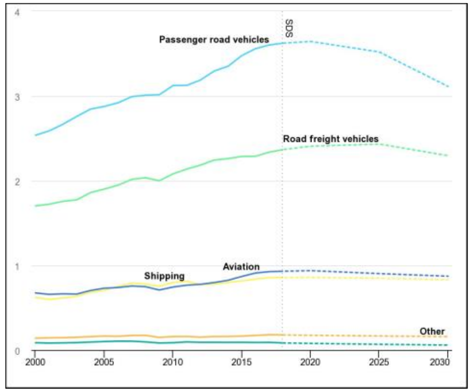

Transport sector CO2 emissions by mode in the Sustainable Development Scenario, 2000 – 2030. Source: [2]

Induction loop within the road registers vehicle type and speed. Source: [6]

References: [1] – Statista 2020, Global parcel shipping volume; [2] – IEA 2019, Transport sector CO2 emissions; [3] – IEA 2021, Global Energy Review: CO2 Emission in 2020; [4] - ITF 2021, Transport Outlook 2021; [5] – Lynn et al. 2018, Decarbonizing intraregional freight systems with a focus on modal shift; [6] – Bundesamt für Strassen ASTRA, Verkehrszähler.

University of St.Gallen

Objectives

Earth-observing satellites, such as the Sentinel-2 satellites from the European Copernicus Program [7] with spectral resolution of 10m per pixel, are able to **detect moving Commercial Vehicles (CVs) on freeways** due to temporal offsets between the different spectral bands leaving a unique signature (see Figure beside). This project examines if we can **estimate commercial vehicle counts from satellite imagery** and approximate ground measurement counts.

Observations of moving CVs (green boxes) in Sentinel-2 satellite imagery. The “rainbow”-effect is caused by a temporal delay between acquiring the individual bands.

For this approach, we use two different models in sequential order. First, we use a **Faster R-CNN** [8] with a **ResNet-50 backbone** as object detector to predict and count commercial vehicles in satellite images. Then, we include **gradient-boosted CatBoost** model [9] to regress the prediction on hourly values to compare it to existing ground-traffic measurements.

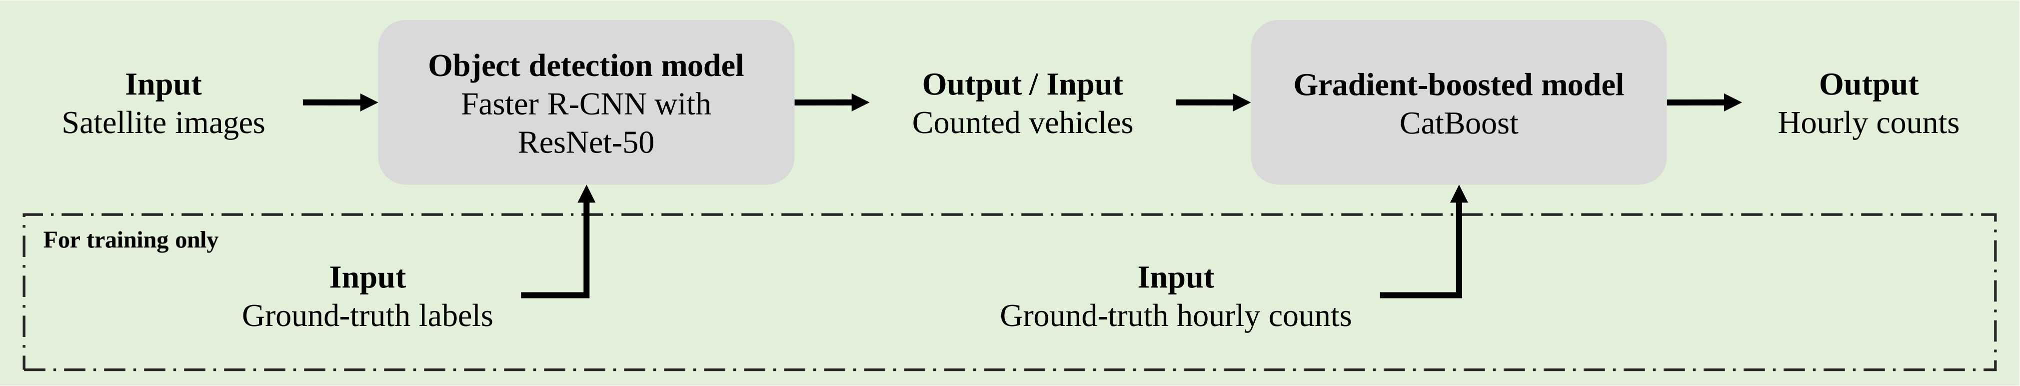

References: [7] – ESA 2015, Copernicus Sentinel-2; [8] – Ren et al. 2015, Faster R-CNN, 1137-1149; [9] – Prokhorenkova et al. 2018, CatBoost.

Data

We have selected several highway sections in Switzerland for our analysis based on 3 criteria:

- Minimum **freeway section length** above 5km
- No **entry & exit points** along the section as they would lead to in- and outflow of vehicles
- **Diverse geographic coverage** of Switzerland.

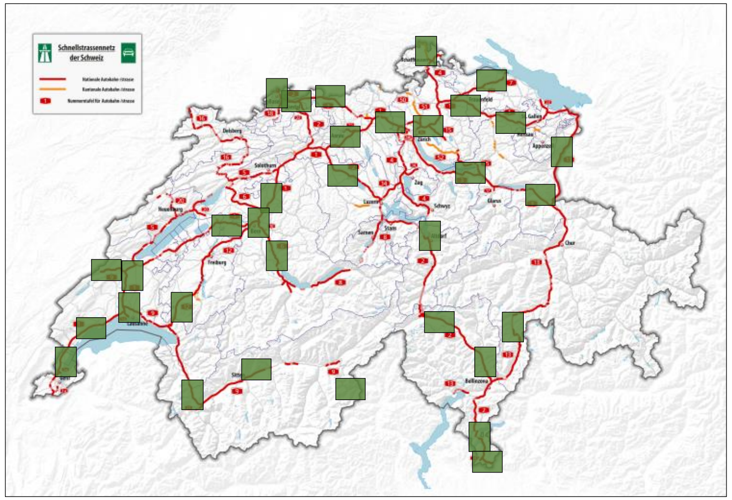

Location of the 33 selected highway sections in Switzerland.

We have collected **satellite imagery** and **ground-traffic measurements** for all selected highway sections for the **year 2019 and 2020.**

Satellite imagery [10]

Collected 3’670 images and created image tiles of 100px times 100px. Annotated 5’683 labeled images with 4’686 individual CVs.

Satellite images from ESA [10].

Ground-traffic data [11]

We gathered hourly ground-traffic measurements from the ‘Bundesamt für Strassen ASTRA Schweiz’.

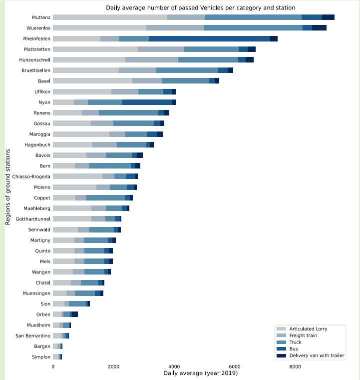

Traffic density across all 33 sections.

References: [10] – ESA Copernicus Open Access Hub; [11] – Bundesamt für Strassen ASTRA 2020, Schweizerische automatische Strassenverkehrszählung.

University of St.Gallen

Detection of Commercial Vehicles

We trained our Faster R-CNN for 400 epochs and reach an **Average Precision (AP) score of 0.72** at an intersection over union (IoU) score of 0.1. The model successfully identifies most of the visible CVs, but occasionally misclassifies other rainbow features or overlooks present CVs. We find, that by **increasing the IoU score to 0.5, the AP drops to 0.65**, as the objects only cover few pixels at once.

Overall, the model tends to make false predictions in cloudy images, as fast-moving clouds generate similar spectral signatures.

MODEL	AP @ IOU 0.1	AP @ IOU 0.3	AP @ IOU 0.5
Faster R-CNN	0.71	0.70	0.64
Faster R-CNN with data augmentation	**0.72**	0.70	**0.65**

Different detection results outcome with green boxes corresponding to ground-truths and red boxes to predictions. With (a) true positives, (b) true negatives, (c) false negative and (d) false positive.

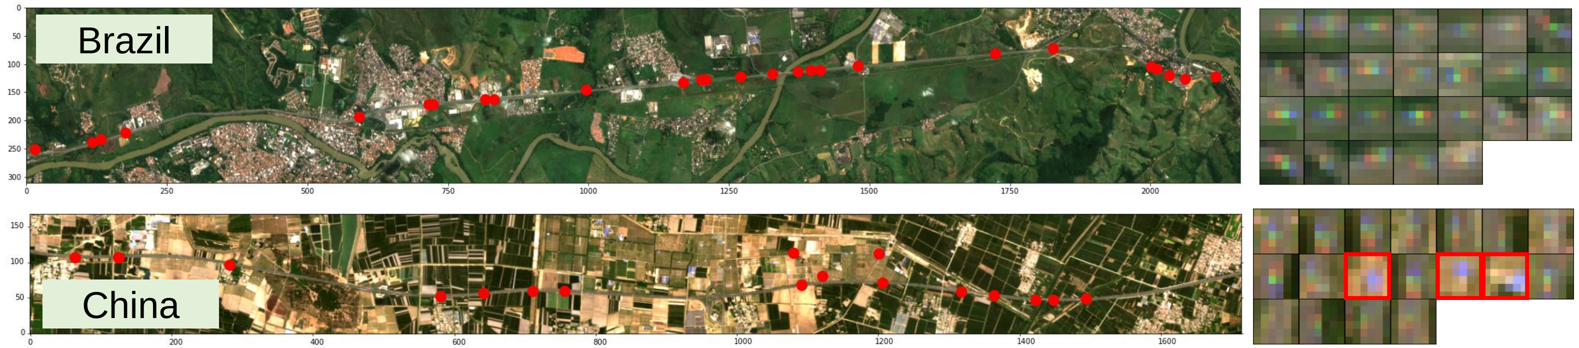

Detection results of our model on highway sections in Brazil and China. Each red dot is a prediction made by our model and is further listed on the right. The detections framed in red refer to predictions besides the highway.

For validation, we also tested our model on five **different geographic areas** with different vegetation and soils. We achieved very similar precision and recall values proofing, that the model is internationally applicable with different sized and looking CVs.

Estimation of Hourly Traffic Counts

As a satellite image is only able to depict a snapshot in time, we utilized a **CatBoost regression model** to map from the number of predicted **CVs per image to hourly counts**.

For the baseline model we extrapolated hourly counts with:

- Number of predictions
- Freeway length
- Speed limit for CV in Switzerland 80km/h

$$\text{Hourly count} = \text{Number of predictions} * \frac{\textit{Freeway length (m)}}{\textit{Speed of trucks}\left(\frac{m}{s}\right)*3600s}$$

Hourly traffic counts baseline model

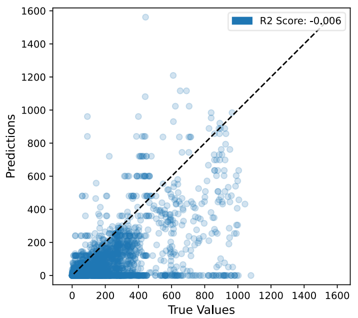

Evaluation Metrics	Baseline
R^2-score	-0.01
Mean Absolute Error	145.98
Root Mean Squared Error	227.87

We further enriched each observation with the **percentage of cloud coverage** across the image as well as the **weekday** during time of capturing. With these features, we trained a **global setting**, i.e., one model for all 33 sections and a **local setting**, i.e., one model per highway section.

Global model

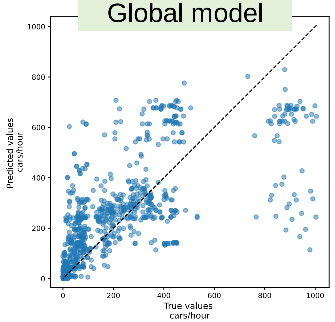

Evaluation Metrics	Global Model
R^2-score	0.46
Mean Absolute Error	122.56
Root Mean Squared Error	180.04

Local model

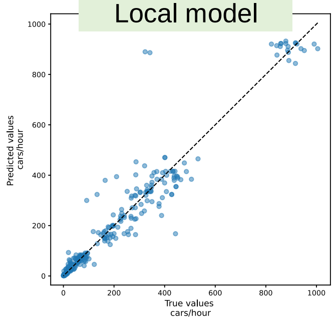

Evaluation Metrics	Local Models
R^2-score	0.93
Mean Absolute Error	33.89
Root Mean Squared Error	66.21

University of St.Gallen

Impact of CoVID-19 Lockdown

We analyzed the influence of the **first wave of the CoVID-19 lockdown** measures on CV traffic rates with a focus on freeway sections in **border proximity**, as they were predominantly affected. The figure below shows the reduction seen over all border areas. We see that the border to Italy (Moleno) was strongly affected, followed by the border to France (Nyon). The border to Germany (Basel) was barely affected by the lockdown measures. Our CV traffic rate predictions **agree well with the ground-truth measurements.**

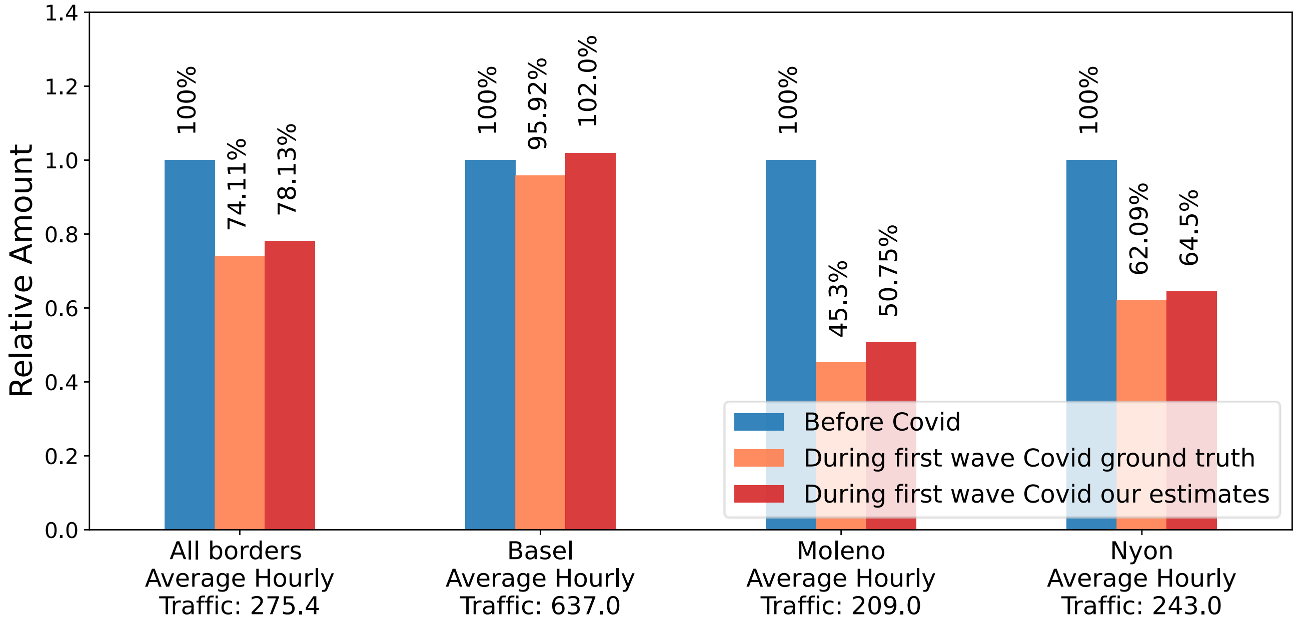

Discussion

Object detection model

We find that the object detection model is rarely confused with false positives, which is further adjustable by **altering the confidence score** of the model. High confidence scores would lead to a conservative model, whereas low confidence scores lead to many false positives. **Masking out the highway** within the satellite images **only leads to a slight performance** improvement, which indicates that the model is very good in learning the context around an annotation. Certain shortcomings can be traced back to the **manual labeling process** of the images. As the objects only cover few pixels, positive annotations could have been overlooked.

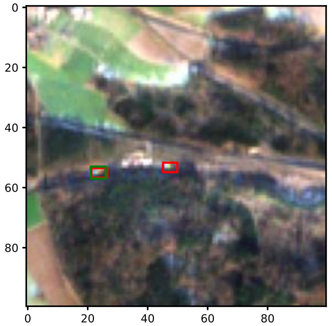

The model correctly identified two CVs within this image, whereas the ground-truth only consisted of one annotation.

Hourly traffic prediction model

By utilizing regression models with observation specific features, we can increase the R^2-score to **0.46 for the global model**, and with historical data available even to **0.93 in the local setting**.

We find, that cloud coverage and highway section length play an important role for the performance of the model. Generally, our model is best applied to observations with a **cloud coverage of less than 20%** as well as on highway sections with a **length above 1km**.

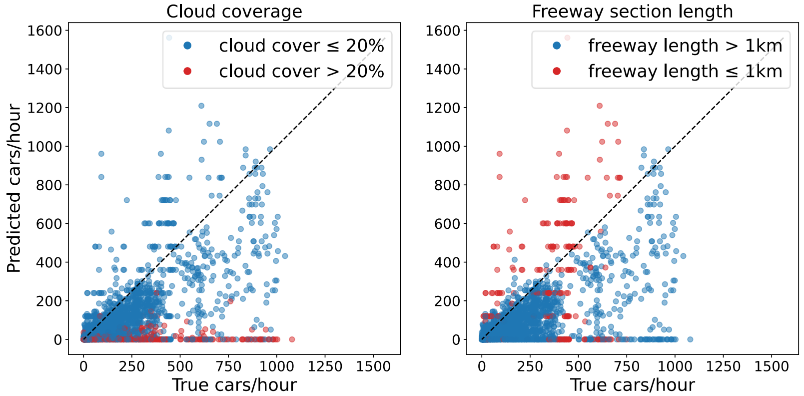

Distribution of observations in the baseline model with cloud coverages below and above 20% as well as highway lengths below and above 1km.

University of St.Gallen

Conclusion

We investigate to possibility to train multiple networks to **detect CVs from satellite images and based of the predictions calculate hourly traffic numbers.** We acquire the satellite data from the **Sentinel-2** constellation from the Copernicus program of the ESA for a total of 33 highway sections within Switzerland. Simultaneously, we collect the ground-traffic measurements from the ‘Bundesamt für Strassen ASTRA Schweiz’.

We are able to achieve an **average precision score of 0.72** for our object detection model in Sentinel-2 data. We find that the model is well able to **learn the context** of true positives and therefore only shows **minimal false positives besides a highway road.**

Based of the predictions we train a CatBoost regression model, that further takes into account the cloud coverage percentage as well as the weekday of the image capturing time and can predict **global traffic measures with a MAPE of 58% and a RMSE of ~160 vehicles per hour**. For freeway sections with **historic data available these metrics drop to 4% for the MAPE and ~60 for the RMSE**.

We conclude that our model pipeline can be utilized to **estimate CV traffic rates on a global scale** to support civil engineers and policymakers in their goal to monitor and minimize greenhouse emissions.

Responsible author:
Moritz Blattner
moritz.blattner@student.unisg.ch