Predicting Traffic Intensity on Motorway Sections

Matic Kladnik[†]
Jozef Stefan International Postgraduate School
Ljubljana, Slovenia
matic.kladnik@gmail.com

Dunja Mladenić
Department of Artificial Intelligence
Jozef Stefan Institute
Ljubljana, Slovenia
dunja.mladenic@ijs.si

Abstract / Povzetek

This paper addresses predictions of traffic intensity on sections of motorways. Predictions are computed for timespans from 24 hours up to 52 weeks. With our adaptive system, we update predictions with newer ones, once additional features can be computed from available data. We use historic context of past traffic intensities on specific sections at specific periods of time, as well as semantic context about the target period. We have evaluated our methodology with multiple machine learning models and compared performances for various timespans on a specific motorway section. The evaluation results show that our methodology improves predictions for specific periods over time.

Keywords

Motorway, traffic intensity, prediction, regression, system, semantic context, evaluation, machine learning

1 INTRODUCTION

A prediction system for predicting traffic intensity on motorway sections can support a wide range of decision making, strategic, and operative processes at the motorway management organization. It can also support end users, such as daily commuters, tourists, and other drivers with their planning of a trip.

The focus of this paper is on architecture of the motorway traffic intensity prediction system as well as on the evaluation of the machine learning models that were trained to produce the predictions for various timespans.

2 PROBLEM SETTING AND DATA

The objective of the proposed methodology is to make long term and medium-term predictions of traffic density or flow (frequency) on various sections of motorway based on historic data of traffic counters, semantic context of motorway stations, semantic context of time periods, and weather data. Predictions serve the motorway management company for better planning of construction projects and to find the least intrusive time slots for road maintenance work. It also serves the motorway drivers when planning a trip.

2.1 Traffic Counters

There are close to one hundred traffic counters that we consider for predictions. Each counter is supported by a pair of inductive loops that are laid into the asphalt of the road. Signals are processed, sent through an IoT communication device and stored into the database.

In the data, there are counts or frequencies of total vehicles, and counts by vehicle types (passenger car, transport truck, bus) for each hour-long time period. E.g. number of vehicles from 8:00 to 9:00 for each of the lanes of a specific motorway section separately.

2.2 Semantic Context

For each of the examples in the dataset we produce semantic context features. For each day and time of day period, we produce semantic context features to inform the model whether a certain time period is on a workday or a weekend, whether the specific time period falls into the morning rush hours or the afternoon rush hours. These semantic features give additional information to improve the performance of machine learning models.

2.3 Data Processing

After downloading the data from the motorway counters via an API of the data provider, we additionally process it to increase consistency and reliability of predictions.

During data processing, we merge data from all lanes of a specific motorway section, which is usually denoted with neighboring towns and the direction of the motorway section.

3 METHODOLOGY DESCRIPTION

We propose a prediction system that includes incorporation of multiple machine learning models to deliver the most reliable predictions based on available data and the timespan for which the system is making predictions of traffic intensity.

To improve prediction accuracy, we make medium-term and long-term predictions. In our case, long-term predictions are made from 1 week to 52 weeks in advance for a specific 1-hour

Information Society 2025, 6–10 October 2025, Ljubljana, Slovenia

http://doi.org/DOI_RECEIVED_AFTER_REVIEW

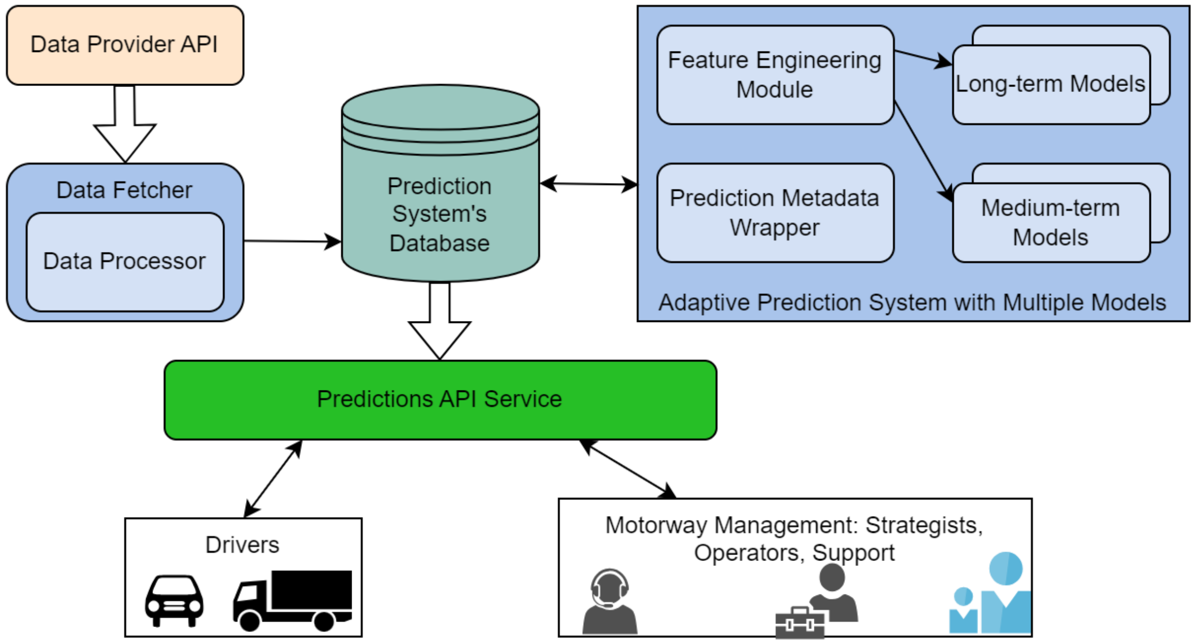

Figure 1: Diagram of the system for producing and distributing predictions of traffic intensity on motorway sections

time period for a specific day of week. Which means that we can make up to 52 predictions when conducting long-term predictions after receiving a new data example, e.g. traffic frequency for a specific 1-hour time period (e.g. 14:00-15:00) for a specific day in time (e.g. Monday).

Whereas medium-term predictions are those that predict from less than 24 hours up to 1 week in advance. For medium-term predictions, we take more features for recent traffic frequency into account for improved accuracy.

Long-term predictions are useful when making decisions for actions that are several weeks or months in the future, while medium-term predictions are more useful when making decisions for actions that will take place from 1 to 7 days in the future.

We have a separate machine learning model for each of the included counters on the motorway to better adjust to specifics of the traffic flow of a specific counter when making predictions of traffic frequency. We have also trained several general-purpose models that are trained on a group of counters or all counters. These are present to support counters with short data history.

Predictions are exposed through a REST API service and are available upon request. They are computed and updated regularly, e.g. daily. More approaches in [1][6].

3.1 Machine Learning Models

To compute predictions of traffic intensity in the future, we use regression machine learning models. We have trained and evaluated several models with the usage of different machine learning algorithms. These are: linear regression, SVM (SVR – Support Vector Machine for Regression), and XGBoost, which is an ensemble model of decision trees.

Features for training models and making predictions are engineered in such a way that each one of the models can use the whole set of features. E.g. we use a one-hot encoding approach when a feature would otherwise have multiple categorical values. We focus on training a specific model for each of the motorway sections that were part of the research. Note that a more general model, trained on data from multiple motorway sections could be more appropriate for motorway sections that have been newly added and do not have enough historical data to support training of a reliable machine learning model with sufficient evaluation period.

Model training processes use MAPE (Mean Absolute Percentage Error, used interchangeably with MARE – Mean Absolute Relative Error). More on relevant machine learning models and metrics in references ([2][3][4][5]).

3.2 Prediction System Description

We continue with the description of our proposed prediction system. The system consists of two main subsystems. One for periodically computing and storing traffic intensity predictions for various time spans. And another for delivering predicted traffic intensity via a REST API service.

As we can see on Figure 1, the system fetches data from the data provider's REST API service. Data is processed after retrieval and sent into a table of prediction system's database. This data is read periodically by the adaptive prediction system.

Once a new value is processed by the system, it checks if there are any additional models with a shorter timespan available, compared to the model used for the currently available prediction. The system prioritizes predictions from models with a shorter timespan in order to update the database with the most reliable predictions available at the time. E.g., prediction with a 1-month timespan succeeds and replaces the prediction with a 3-month timespan.

Different long-term and medium-term models can be trained using different machine learning algorithms, depending on the algorithm that performed the best during the evaluation of the models.

Once updated the predictions are stored in the database, they are available to users, such as strategists, operators and support specialists within the motorway management organization. Or end users of the motorway, such as drivers of cars, trucks, buses, etc. A key advantage of this approach is that drivers and motorway operators and specialists get insights that are based on the same predictions for traffic intensity, which supports greater transparency of information and stronger compatibility of different applications for end users and motorway professionals.

E.g. the system can support long-term planning for larger maintenance or reconstruction projects for up to 1 year ahead, as well as long-term planning of road users. For instance, drivers can plan their holidays and the time of their commute ahead. And highway maintenance operators can find the most optimal schedule for short maintenance work.

4 EVALUATION

We continue with the evaluation of the machine learning models. To compare models, trained with different algorithms, we use the evaluation results for the same motorway section on the Slovenian motorways. We use the period from 1 May 2024 until 5 May 2025 for evaluation.

We use Scikit-learn library[7] to train the linear regression (using ordinary least squares approach) and SVM (SVR) models and the XGBoost library[8] to train the XGBoost models. SVM model is trained using the RBF kernel, and with scaled gamma hyperparameter. In majority of motorway sections, XGBoost models with a maximum depth of 6 performed the best which is why we used models with the same hyperparameter value for the following analyses. We use gbtree as the booster, while the learning rate is 0.3.

Table 1: Model Performance Comparison

timespan	algorithm	MAE	RMSE	MAPE
24 hours	XGB	39.43	62.75	10.5%
24 hours	SVM	42.38	65.86	11.5%
24 hours	lin. reg.	43.14	66.93	11.6%
7 days	XGB	45.66	70.69	11.6%
7 days	SVM	43.70	68.91	12.1%
7 days	lin. reg.	43.51	69.04	12.1%
4 weeks	XGB	57.30	88.56	13.9%
4 weeks	SVM	50.20	77.86	14.1%
4 weeks	lin. reg.	51.33	78.63	14.7%
52 weeks	XGB	88.33	121.93	20.9%
52 weeks	SVM	53.54	84.49	14.9%
52 weeks	lin. reg.	70.46	96.98	21.3%

We evaluated the models on a little over 1 year of test data, which was not included in the training or validation part of the process.

We continue with the analysis of the model performances as seen in Table 1. If the timespan attribute's value is '7 days', it means that the model predicts 7 days into the future. We use several metrics to describe the performance of the models. These are: MAE (Mean Absolute Value), RMSE (Root Mean Square Error), and MAPE (Mean Absolute Percentage Error). MAPE is a crucial metric as it shows relative errors in percentages which is key when evaluating the models as traffic frequency varies significantly throughout different parts of the day.

We can see some interesting performance dynamics of the models. The XGBoost model performs the best for 24-hour timespan, with a significant performance uplift of at least 1 percentage point in MAPE, compared to the other two models. It is also better in the other two metrics: MAE and RMSE. We consider predictions with 24-hour timespan to be medium-term predictions.

We continue with the performance analysis of the long-term predictions. For the 7-day timespan, the XGBoost model is still noticeably better than the other two models with a 0.5 percentage point uplift in performance. For the 4-week timespan, XGBoost still holds a small lead in the key metric (MAPE), whereas the SVM model has significantly better results when considering just MAE and RMSE metrics. For the 52-week timespan, we can see an interesting dynamic as the SVM model takes a significant lead in performance as it is the only one with the MAPE value of less than 15%, whereas the MAPE values of the other two models surpass 20%.

The dynamic is likely caused by a reduced set of features as there are significantly less historic traffic count features that are included when making predictions with a 52- week timespan. It seems this has a significantly negative impact on training the XGBoost model, which is a tree ensemble model, while having additional features available gave the XGBoost model an edge for predictions with a timespan up to 4 weeks, especially up to 7 days.

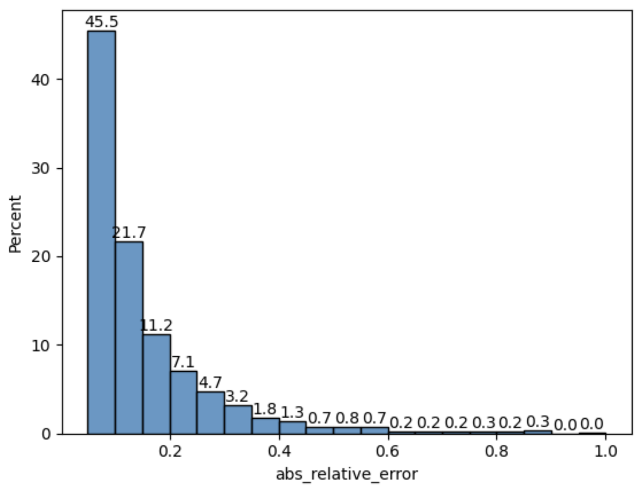

Figure 2: Distribution of absolute relative errors by 5% buckets for XGBoost 7-day timespan model

On Figure 2 we can see how absolute relative errors are distributed if they are split into 5% absolute relative error buckets. We can see that in 45.5% of the cases, the absolute relative (or percentage) error of the predicted traffic frequency is less than 5% of the actually measured traffic frequency. 21.7%

of predictions have a relative error between at least 5 and (excluding) 10 percent, and 11.2% of predictions have a relative error between 10 and 15 percent.

This means that in 78.4% of predictions, the relative error was less than 15%, which can be considered as a sufficiently good performance for the models to support a sufficiently reliable traffic intensity prediction system.

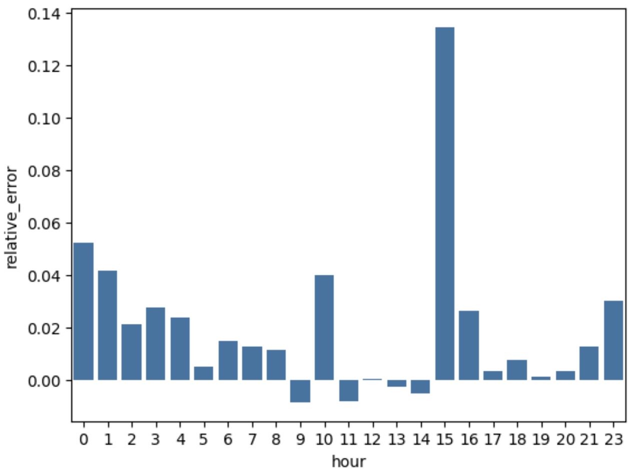

Figure 3: Mean relative errors by each hour of the day for XGBoost 7-day timespan model

We continue by analyzing the distribution of mean relative errors by each hour of the day as seen on Figure 3. We can see that the model generally tends to slightly overestimate or overshoot with its predictions. Especially during the night-time periods, when there are fewer vehicles on the motorway.

In the mean aggregate, there is less than a 2% mean relative error during the morning rush hours (at 6:00-7:00, 7:00-8:00, and 8:00-9:00). It is the highest during the 15:00-16:00 period, with more than 13% of mean relative error. However, the error is substantially smaller during other afternoon rush-hour periods, 14:00-15:00, 16:00-17:00, and 17:00-18:00, where it remains under 4%. Apart from the 15:00-16:00 period, the mean relative errors are consistently under 6%. When the model does undershoot or underestimate with its prediction, the mean relative error is less than 2%, close to 1%.

We can see a spike of mean relative error at the 15:00-16:00 period. Upon investigation, it seems likely it was caused by an issue somewhere in the process of data collection as there was an irregularly high mean relative error for the month of May.

4.1 Evaluation Insights

When considering the results of the evaluation of trained machine learning models for specific motorway sections, we have gathered several key insights.

In some examples, we could not compute all features due to missing values in data, meaning that certain features had NaN values after computing historic time-series features with Pandas' shift function. In this case there is a strong advantage of having a decision tree ensemble model (e.g. XGBoost) as a backup, even if it is not the best performing model for a certain timespan. This is due to the ability of the tree ensemble models to apply only those trees that are covered by features with available values. In this case the predictions are generally less accurate. However, we can still compute them.

Another key insight is that the evaluation supports our proposed methodology with multiple models to improve the performance of the predictions for each included timespan.

Another useful insight is that different algorithms can produce the best models for different timespans on the same motorway section. As was the case with the SVM model in our evaluation.

5 CONCLUSION

We have overviewed the methodology that we use as the foundation for our proposed system for predicting traffic intensities on motorway sections. Including the adaptive prediction system and the supporting machine learning models that support making predictions for various timespans to, in time, improve already available predictions for specific time periods in the future. We have also overviewed the evaluation of the trained machine learning models and found some useful insights that support our proposed prediction system.

Based on the current evaluation results that were presented in the paper, our methodology produces predictions with sufficient reliability to support long-term decision making of various roles.

For further improvements to the system, we could train and evaluate some deep learning models and models that are based on the transformer architecture, as well as some other time-series forecasting procedures, such as Facebook Prophet. We could also engineer additional semantic context features for further improvements to the performance of the existing models. For additional improvements for shorter timespans, we could also include weather forecast data.

References / Literatura

[1] Bernardo Gomes, Jose Coelho, Helena Aidos. 2023. A survey on traffic flow prediction and classification. In *Intelligent Systems with Applications*, vol. 20. DOI: https://doi.org/10.1016/j.iswa.2023.200268

[2] Jithin Raj, Hareesh Bahuleyan, Lelitha Devi Vanajakshi. 2016. Application of Data Mining Techniques for Traffic Density Estimation and Prediction. *Transportation Research Procedia, vol 17*. DOI: https://doi.org/10.1016/j.trpro.2016.11.102

[3] Yuyu Zhu, QingE Wu, Na Xiao. 2022. Research on highway traffic flow prediction model and decision-making method. *Scientific Reports*, vol. 12. DOI: https://doi.org/10.1038/s41598-022-24469-y

[4] Carl Goves, Robin North, Ryan Johnston, Graham Fletcher. 2016. Short Term Traffic Prediction on the UK Motorway Network Using Neural Networks. *Transportation Research Procedia, vol. 13, 184-195.* DOI: https://doi.org/10.1016/j.trpro.2016.05.019

[5] Adriana-Simona Mihaita; Zac Papachatgis; Marian-Andrei Rizoiu. 2020. Graph modelling approaches for motorway traffic flow prediction. 2020. *IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. DOI: https://doi.org/10.1109/ITSC45102.2020.9294744

[6] Sayed A. Sayed, Yasser Abdel-Hamid, and Hesham A. Hefny, 2022. Artificial Intelligence-Based Traffic Flow Prediction: A Comprehensive Review. *Pre-review*. DOI: http://dx.doi.org/10.21203/rs.3.rs-1885747/v1.

[7] Scikit-learn: https://scikit-learn.org

[8] XGBoost: https://xgboost.ai/