The effect of spatial-temporal dependencies on forecasting traffic sensor detections

BEUTH HOCHSCHULE FÜR TECHNIK BERLIN
University of Applied Sciences

Felix Kunde, Alex Rudi, Anatoli Procik, René Jablonski, Tatjana Manych, Alexander Hartenstein, Petra Sauer

With decreasing costs for sensors, capturing spatial-temporal phenomena in real time becomes feasible and yet more powerful. A greater ground truth helps to build better predictive models that can be used to create superior planning strategies. Different algorithms are suitable for time series analysis on sensor data. Due to the current hype around machine learning recent efforts cycle mostly around the idea to adopt various kinds of artificial neural networks (ANN). Within the ExCELL research project we employed different ANNs on double loop detector data with generally good results even without a lot of fine tuning. We wanted to test the generalizability of our models on a different dataset with more irregular traffic patterns and chose a freeway scenario in Los Angeles with data coming from the PeMS portal by the State of California.

Keywords: Traffic Prediction, Time Series, LSTM Neural Network, Spatial-temporal Data Mining

Positive and negative effects:

+ stable traffic pattern
+ expressive input features
+ neighbours with similar patterns
- volatile input features
- neighbours with high prediction errors
- spatial weight matrix introduces a bias

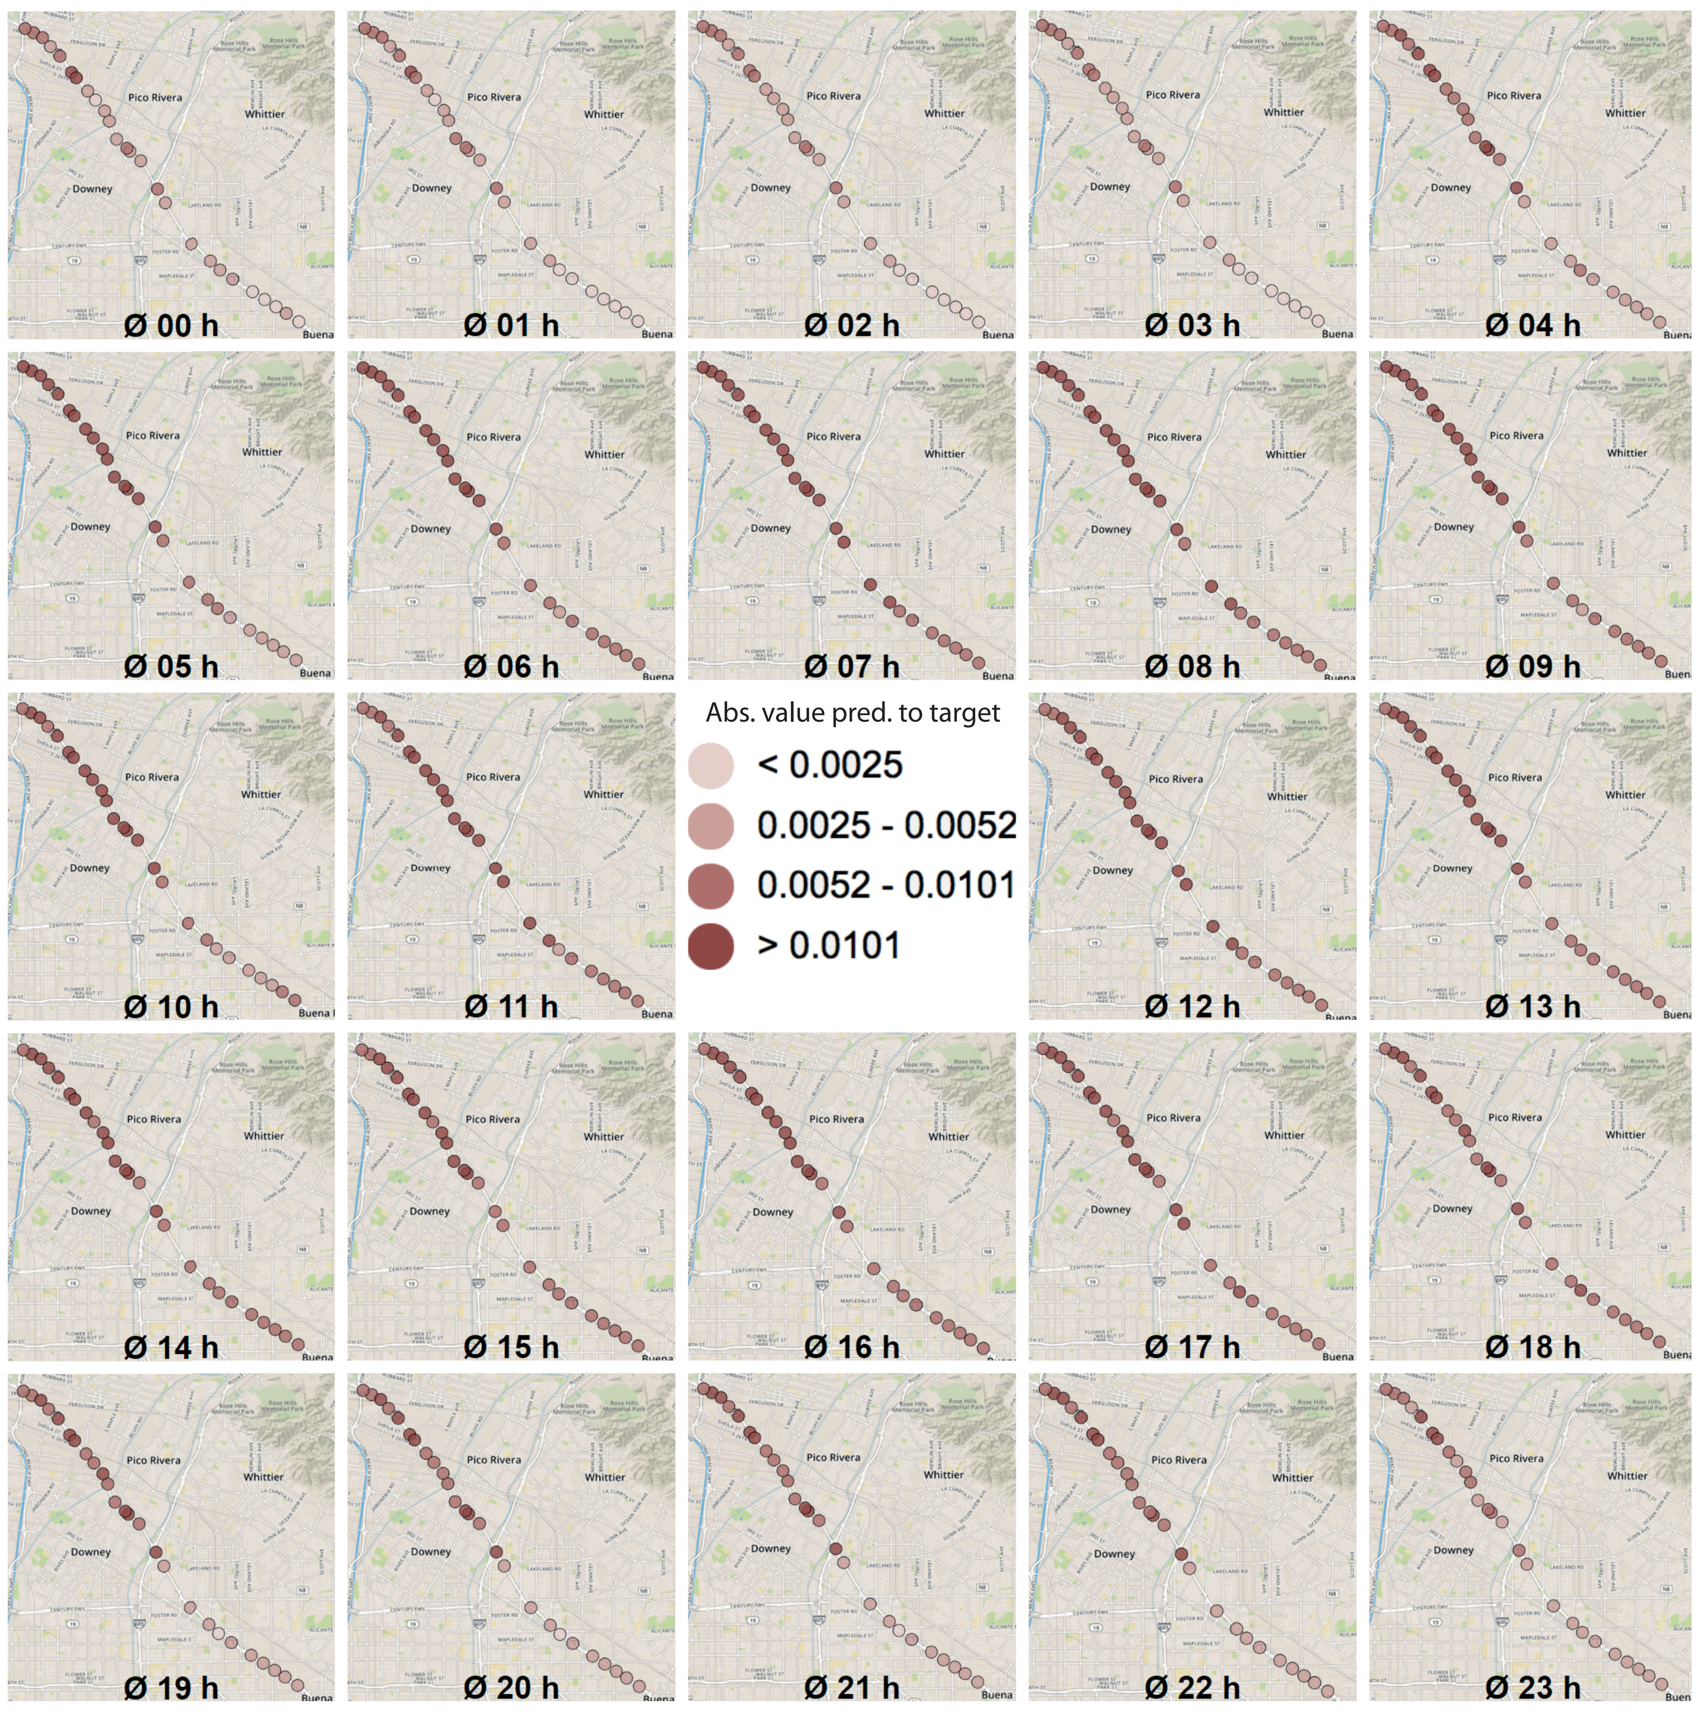

Figure 1: Average deviation from 5 minutes predictions of occupancy values per hour of day for one year

Figure 2: Average Level of Service - LOS (vehicles/mile/lane) per hour of day for one year

Time Series Analysis with LSTMs

Our prediction model uses the Long Short-Term Memory ANN (LSTM) implementation in Tensor Flow. LSTMs are designed to learn on sequence-based data and are heavily used in text and speech recognition. But, during the last years they have also been tested in various other domains outperforming regression-based machine learning algorithms (Greff et al. 2017). Our implementation can be found under our GitHub account: github.com/MAGDa-BeuthHS/dlsd.

Training Data

We took data of one year from 27 inductive loop sensors of highway no. 5 covering a section from East Los Angeles till Buena Park (see maps in figure 1 and 2). The loops detect different parameters relevant for traffic monitoring and planning such as flow (number of cars), occupancy (time on sensor) and speed aggregated by an interval of 5 minutes. Gaps of missing data are filled up with linear interpolation in R. A matrix of timestamps and normalized values is used to train the model. We use k-fold cross validation to split the data into training and test data sets. Grid search has been used to find a good hyperparmeter setup for the LSTM. The overall prediction accuracy was best when using occupancy as the input variable. Probably because it is the most adequate measure for the current traffic situation, a.k.a. the Level of Service (HCM 2000).

Research Goal

The trained model predicts eight different future values for each sensor ranging from 5 to 90 minutes. The main goal of this research is to tell where and when the predictions work better, where and when worse. By simply plotting a distribution of the speed values at each sensor we get a first idea, which parts of the highway suffer more from congestion than others (see figure 3).

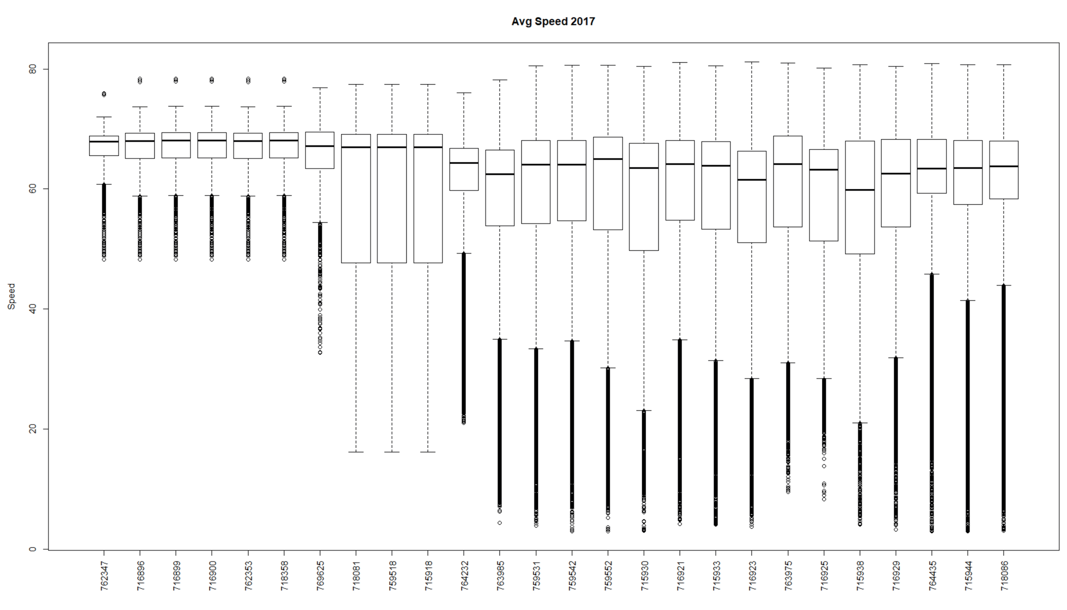

Figure 3: Average speed distribution on all sensors for one year

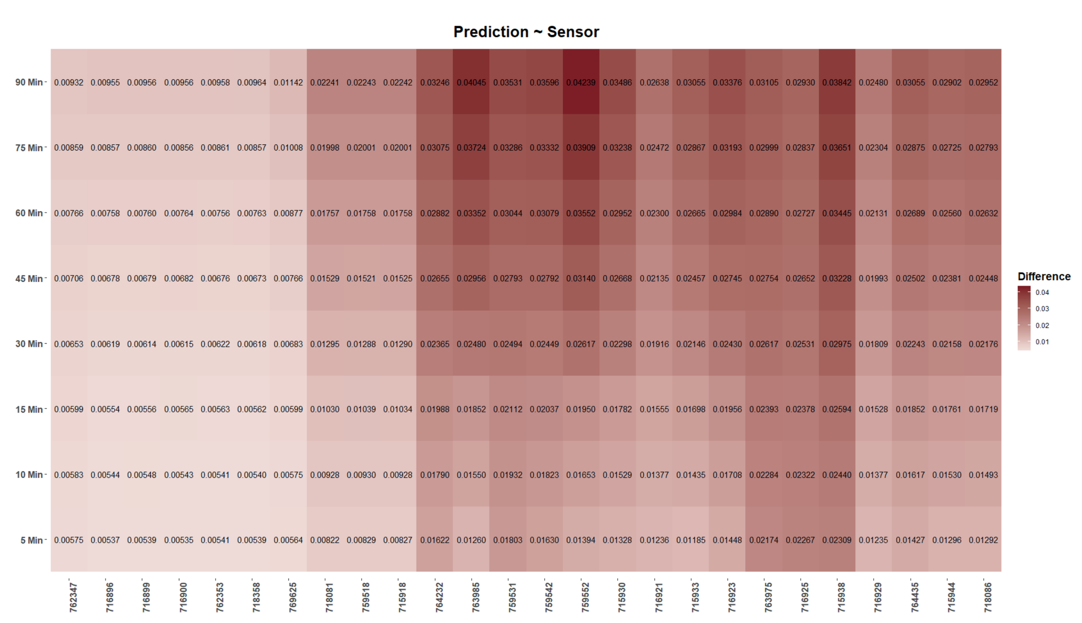

Figure 4: Average diffenrence to predictions of all sensors for one year

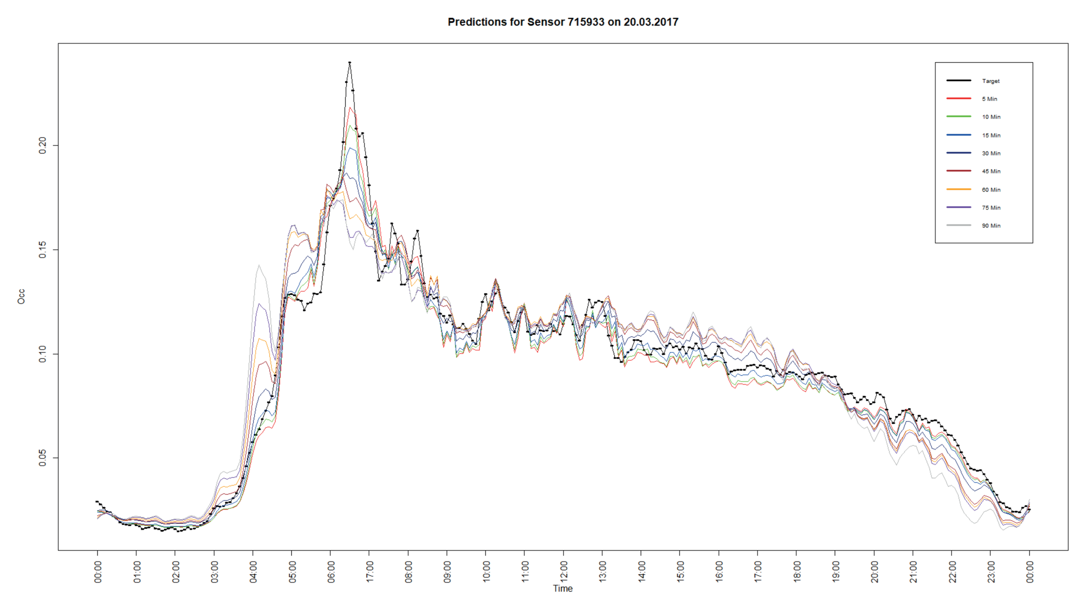

Figure 5: Prediction frames compared to real occupancy for one sensor on a monday

Results

From the figures we can conclude that forecasting works better when less congestion occurs, or more precisely, when the LOS stays stable at one level. For the first six sensors the model returns accurate predictions even for a 90 minutes time frame. The highest errors occur at sensors with a bad LOS (D-E). In figure 5 the deviation to the target value is plotted for one sensor which shows that detection sequences during peaks are not learned well for long-term predictions. The increase and decrease of road occupancy are predicted too early. The higher the prediction time frame the worse the error.

A similar finding has been made in a previous study on arterial road network data of Dresden, Germany (Kunde et al. 2017). Like before, we tried to include a spatial weight matrix in the training but the predictions got even worse for the freeway setting. For future research we plan to include more features and add a convolutional layer to improve the filtering of spatial relationships between sensors.

FOSS4G 2018 DAR ES SALAAM

MAGDa Management & Analytics of Geotemporal Data

excell
excell-mobility.de/developer

DLR

Supported by:

Federal Ministry for Economic Affairs and Energy

on the basis of a decision by the German Bundestag

References:

Greff, K.; Srivastava, R.K.; Koutnik, J.; Steunebrink, B.R.; Schmidhuber, J.: LSTM: A search space odyssey. IEEE transactions on neural networks and leraning systems 28(10) (2017).

Kunde, F.; Hartenstein, A.; Sauer, P.: Spatial-temporal Traffic Flow Forecasting on a City-wide Sensor Network. GIS Ostrava 2017 Proceedings.

Transportation Research Board: Highway Capacity Manual 2000. Washinghton D.C.