Supplemental Material : Dynamic Traffic Modeling from Overhead Imagery

Scott Workman
DZYNE Technologies

Nathan Jacobs
University of Kentucky

This document contains additional details and experiments related to our methods.

1. Dynamic Traffic Speeds Dataset

We presented a new dataset for fine-grained road understanding containing 11 902 non-overlapping overhead images, associated road attributes, and historical traffic data. Figure 1 shows the spatial coverage of our dataset, with yellow, blue, and magenta corresponding to the location of training, testing, and validation images, respectively. Figure 3 shows example images from the dataset along with the road mask and a speed mask rendered using a random time (road segments buffered to two meter half width).

We also report some statistics of the underlying historical traffic speed data. Note that these numbers are computed on the provided aggregated traffic speed data, which is then further aggregated by day of week and hour of day. The average road segment speed in NYC for 2018 (averaging road segments over time first, then averaging across segments) is approximately 30.47 km/h ($\sigma = 11.99$). In Figure 2, we visualize the average road segment speed versus time. Additionally, Figure 4 shows the average speed according to OpenStreetMap's road type classification.

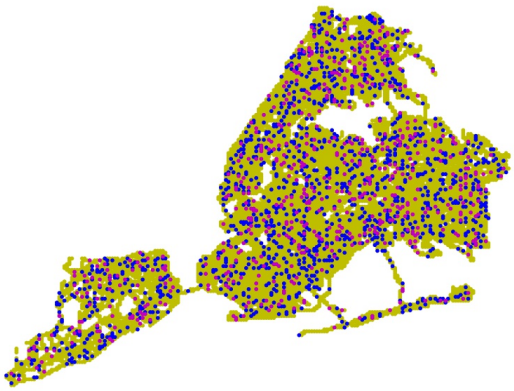

Figure 1: Coverage of our dataset. Each dot corresponds to a non-overlapping overhead image (yellow training, blue testing, magenta validation).

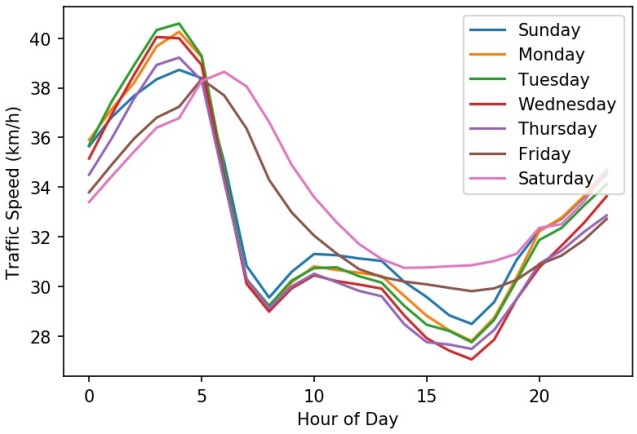

Figure 2: Average road segment speed versus time.

2. Extended Evaluation

In the main document, we presented an evaluation for traffic speed estimation where each image in the test set was associated with a random time to represent the ground truth speed data. These image/time pairs were then fixed for all methods evaluated. Here we explore a macro evaluation, where we consider the relationship between performance and time. Table 1 shows the results of this experiment for a subset of times. When computing metrics, we select all images in the test set which have at least one segment with empirical traffic speed data at the time of interest. Then, we compute metrics for each time and average the result (treating all times equally). For this experiment, we consider Monday and Saturday and selected the following hours of the day: 12am, 4am, 8am, 12pm, 5pm, and 8pm. Consistent with our earlier results, our method that integrates location and time outperforms an image-only baseline.

2.1. Impact of Angle-Dependent Speeds

In our approach, we estimate angle-dependent speeds as opposed to predicting a single speed per location. The intu-

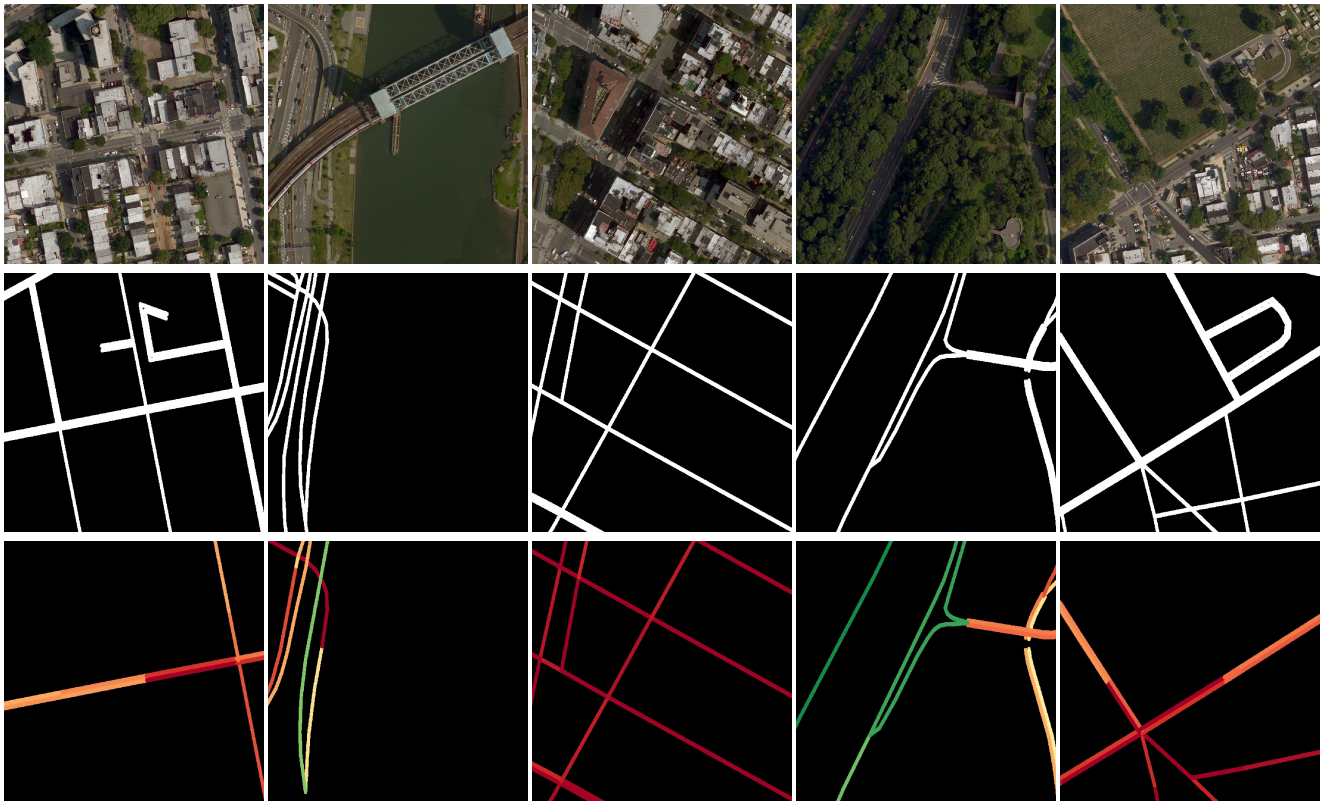

Figure 3: Examples from our dataset: (top) image, (middle) road mask, and (bottom) speed mask rendered using a random time (where red is slower and green faster). Notice that historical speed data is not available for every road at every time.

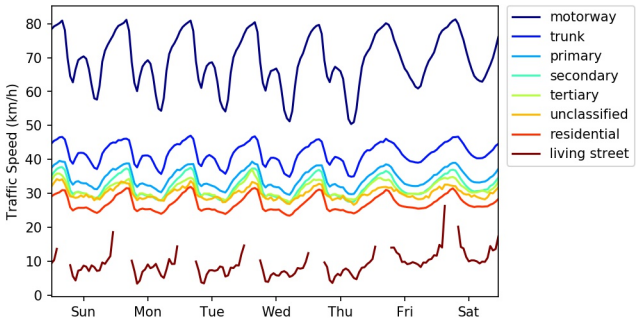

Figure 4: Average road segment speed versus time, where each road is categorized by its OpenStreetMap road type classification.

ition behind this idea is that speed tends to depend on direction, e.g., a bridge that crosses over a highway. Following the above evaluation scheme, we compare our strategy of making angle-dependent speed predictions to a baseline that instead uses uniform weights. In other words, we replace the orientation-weighted average with an equal-weighted average. The results are shown in Table 1. Our full approach outperforms the uniform variant.

Table 1: Quantitative evaluation of traffic speed estimation (RMSE).

	image	*Ours (uniform)*	*Ours*
Monday (4am)	15.82	13.31	**13.24**
Monday (12pm)	10.96	10.59	**10.41**
Saturday (5pm)	10.65	10.39	**10.36**
Saturday (8pm)	10.43	10.32	**10.27**
Overall	11.903	11.145	**11.134**

3. Application: Augmenting Routing Engines

Our method can be applied to generate optimal travel routes that take into account traffic speeds at different times. For this experiment, we use the OSMnx library [1] to represent the underlying road topology. We compute the traversal time of each edge based on the road segment length and our speed estimate. Figure 5 shows the results of this experiment for a route in Queens, New York. Figure 5 (left) shows the route corresponding to shortest overall distance. Figure 5 (middle) shows the route corresponding to the shortest travel time on Monday at 4am. Figure 5 (right) shows the

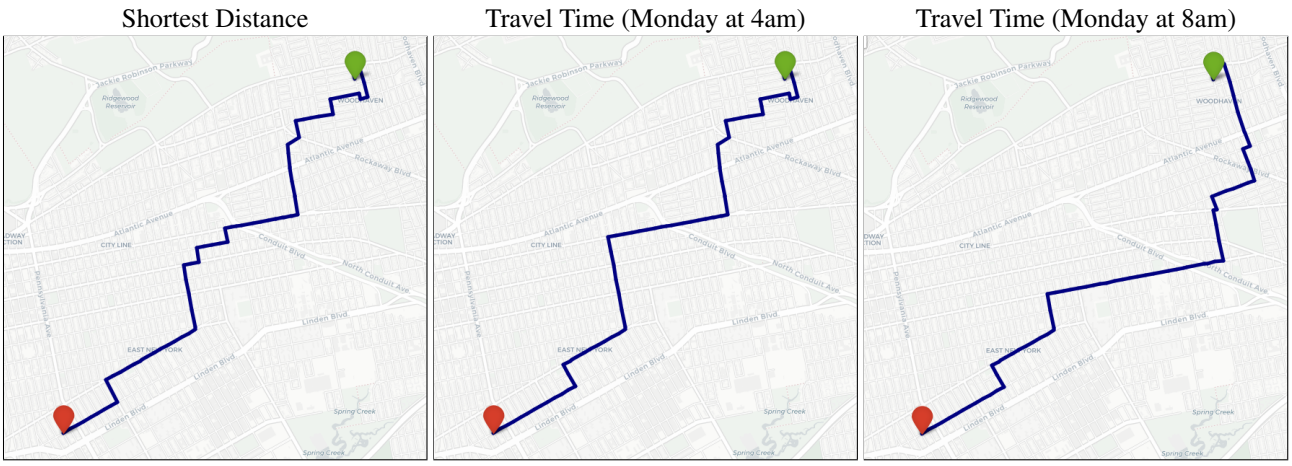

Figure 5: Using dynamic traffic speed predictions to augment route generation. The routes correspond to (left) the shortest path in terms of total length, (middle) travel time on a Monday at 4am, and (right) travel time on a Monday at 8am. Note that for the travel time routes, edge weights are represented by traversal times and computed using the length of the road segment and corresponding traffic speed estimated by our approach.

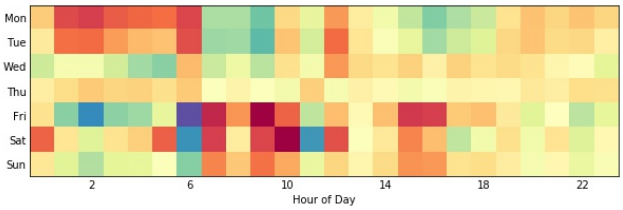

Figure 6: Visualizing the learned time embedding.

route corresponding to the shortest travel time on Monday at 8am.

4. Additional Results

Visualizing the Time Embedding In Figure 6 we visualize the learned time embedding. To form this image, we take the three dimensional embedding for each dimension of time (day of shape 7×3 and hour of shape 24×3) and form a false color image of shape $7 \times 24 \times 3$ by broadcasting; then average across channels. As observed, the learned embeddings are clearly capturing traffic speed patterns related to time. For example, on Friday and Saturday the embedding reflects much larger values (red) around the middle of the day as opposed to other days of the week. This makes sense as many people work half days on Friday and leave work early or take the family shopping on Saturday. This result also agrees with the temporal patterns for traffic speed shown in Figure 2.

Qualitative Results Figure 7 shows several example images alongside the ground-truth road mask and the prediction from our approach. Similarly, Figure 8 shows additional examples for orientation estimation, visualized as a flow field for a subset of points.

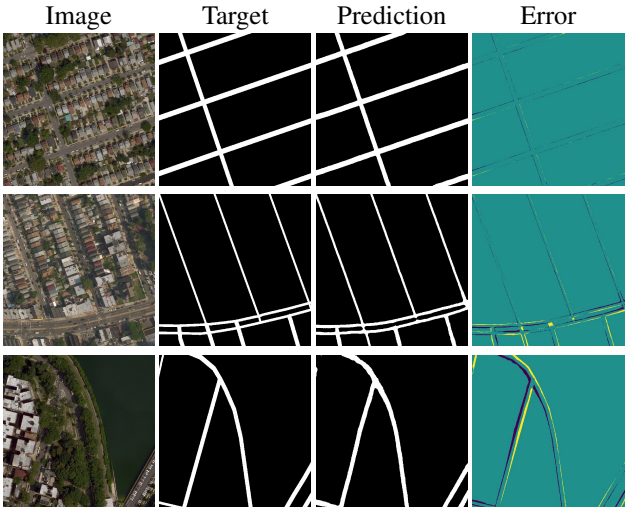

Figure 7: Qualitative results from road segmentation. The error image (right) shows false positives (negatives) color coded as purple (yellow).

References

[1] Geoff Boeing. OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139, 2017. 2

Figure 8: Using our approach to estimate directions of travel (visualized as flow fields).