

# RasterNet: Modeling Free-Flow Speed using LiDAR and Overhead Imagery

Armin Hadzic<sup>1</sup> Hunter Blanton<sup>1</sup> Weilian Song<sup>2</sup> Mei Chen<sup>1</sup> Scott Workman<sup>3</sup> Nathan Jacobs<sup>1</sup>  
<sup>1</sup>University of Kentucky <sup>2</sup>Simon Fraser University <sup>3</sup>DZYNE Technologies

## Abstract

Roadway free-flow speed captures the typical vehicle speed in low traffic conditions. Modeling free-flow speed is an important problem in transportation engineering with applications to a variety of design, operation, planning, and policy decisions of highway systems. Unfortunately, collecting large-scale historical traffic speed data is expensive and time consuming. Traditional approaches for estimating free-flow speed use geometric properties of the underlying road segment, such as grade, curvature, lane width, lateral clearance and access point density, but for many roads such features are unavailable. We propose a fully automated approach, RasterNet, for estimating free-flow speed without the need for explicit geometric features. RasterNet is a neural network that fuses large-scale overhead imagery and aerial LiDAR point clouds using a geospatially consistent raster structure. To support training and evaluation, we introduce a novel dataset combining free-flow speeds of road segments, overhead imagery, and LiDAR point clouds across the state of Kentucky. Our method achieves state-of-the-art results on a benchmark dataset.

## 1. Introduction

Free-flow speed is defined as the average speed a motorist would travel on a given road segment when it is not impeded by other vehicles. This is an important measure used in transportation engineering for a variety of applications such as traffic control, highway design, measuring travel delay, and setting speed limits. Existing approaches for collecting measurements of free-flow speed have largely been manually intensive and difficult to scale [3], putting a large strain on transportation engineering budgets. Only recently have more advanced techniques, such as probe vehicles, been used for road performance monitoring [19]. To avoid the upfront cost of collecting traffic speed data, a variety of recent work has explored developing automatic methods for estimating free-flow speeds.

Traditional approaches for free-flow speed modeling involve the use of geometric road features (also known as highway geometric features) such as lane width, lateral

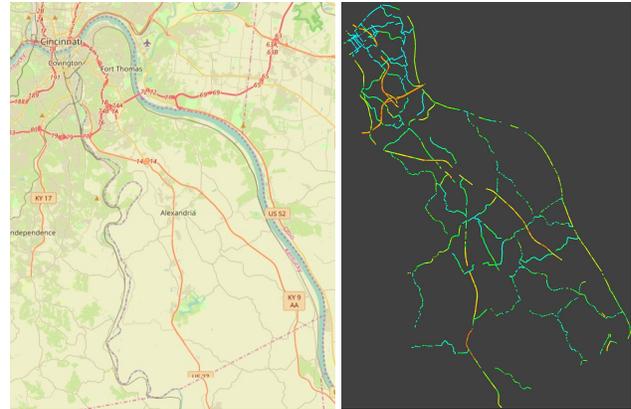


Figure 1: We propose an automatic approach for estimating free-flow speed from overhead imagery and 3D airborne LiDAR data. (left) A map representing Campbell county in Kentucky, USA. (right) The corresponding map of free-flow speeds generated using our method.

clearance, median type, and access points [13]. These approaches tend to be specific to certain road network types (arterial, local, collector) [22], or geographical areas (urban and rural) [20]. While these methods have demonstrated good performance, their use is limited to areas where the necessary road metadata is available. Typically, these areas include state-maintained highways such as interstates, US highways, and state roads. However, this is often a small portion of all roads. For example, only 35% of all roadway miles in Kentucky are state-maintained. The detailed geometric features required for estimating free-flow speed on locally maintained roads are mostly unavailable or prohibitively expensive to collect. Estimating free-flow speeds at large scales requires learning-based methods that take advantage of alternative data sources (Figure 1).

Recent work has shown that road geometry approaches can be augmented with visual data, in the form of overhead imagery, to improve performance [23]. Though adding visual features results in better performance than road geometric features alone, model applicability is still limited to sufficiently documented roads. Instead, we explore replacing explicit road geometric features with features extracted

from airborne LiDAR (Light Detection and Ranging) point clouds. Compared to image data which is often impacted by transient effects (e.g., weather), 3D point clouds are viewpoint invariant, robust to weather and lighting conditions, and provide explicit 3D information not present in 2D imagery, offering a supplementary source of data. Our approach combines both sources, visual features extracted from overhead imagery and geometric features extracted from point clouds.

We propose *RasterNet*, a multi-modal neural network architecture that combines overhead imagery and airborne LiDAR point clouds for the task of free-flow speed estimation. To align the input domains, *RasterNet* organizes local point cloud neighborhoods using a raster center grid and pairs them with spatially consistent features extracted from the image data. Features from both domains are then merged together and used to jointly estimate free-flow speed. To support the training and evaluation of our methods, we introduce a large dataset containing free-flow traffic speeds, overhead imagery, and airborne LiDAR data across the state of Kentucky. We evaluate our method both qualitatively and quantitatively, achieving state-of-the-art results compared to existing methods, without requiring explicit geometric features as input.

Our primary contributions can be summarized as follows:

- A large dataset for free-flow speed estimation that combines speed data, overhead imagery, and corresponding point clouds.
- A novel multi-modal neural network architecture for free-flow speed estimation that advances the state-of-the-art on an existing benchmark dataset.
- A method for fusing overhead imagery and airborne LiDAR point clouds using a geospatially consistent raster structure.

## 2. Related Work

We provide an overview of work in three related fields: point cloud representations multi-modal data fusion, and estimating traffic speed.

### 2.1. Point Cloud Representations

Many methods have been proposed for extracting feature representations from point clouds. Recently, Weinmann et al. [25], Liu et al. [12], and Dubé et al. [5] demonstrated that point clouds could be represented by neighborhood structural statistics in order to improve performance on scene understanding and place recognition tasks. The seminal work of Qi et al. [17] introduced PointNet, a general deep neural network for point cloud feature extraction. This work

inspired a series of works in point cloud shape classification [10, 14, 24] and object detection [21]. Later, Qi et al. presented PointNet++ [18], a shape classification method and extension to PointNet which adds local feature extraction to improve performance. This method allows for precise control over the spatial location of extracted features, which we use for geospatially aligning point cloud features with visual features from an image.

### 2.2. Multi-Modal Data Fusion

A significant amount of work has explored combining imagery with LiDAR data for various tasks. Liang et al. [11] designed a method for multi-scale fusion of ground imagery with overhead LiDAR point clouds to perform object detection from multiple viewpoints and modalities. Similar to our own work, Jaritz et al. [8] used a cross-modal autonomous driving dataset to perform unsupervised domain adaption for 3D semantic Segmentation. Their dataset combined terrestrial LiDAR point clouds and camera images for different times of day, countries, and sensor setups. Their proposed cross-modal model, xMUDA, performs data fusion by projecting 3D point cloud points onto the 2D image plane and sampling features at corresponding pixel locations. While this dataset and method were designed for small spatial areas around a vehicle, we perform data fusion of overhead imagery and airborne LiDAR point clouds of large  $200 \times 200m^2$  areas.

Recent work has also explored the fusion of airborne LiDAR with overhead imagery for the task of semantic segmentation in an urban area [2]. Typically these approaches render the LiDAR data as 2D images through digital surface models and use a traditional CNN. This strategy results in a loss of precise 3D information due to discretization. This is an issue, as raw point cloud methods have been shown to outperform discretization-based approaches for classification tasks [15]. Our approach uses point cloud understanding to process 3D point clouds.

### 2.3. Estimating Traffic Speed

Several works have proposed automatic methods for estimating the speed of vehicles. Huang [7] used video surveillance data of traffic to perform individual vehicle speed estimation. We perform average free-flow speed estimation to characterize traffic flow behavior and capacity of roads instead of individual vehicle speed characteristics. Most similar to our own work, Song et al. [23] performed free-flow speed estimation using overhead imagery and geometric road features on the Kentucky free-flow speed dataset. Our *RasterNet* model is trained on the same overhead imagery and label data, but our approach replaces the provided geometric road features with point cloud features of the same spatial area.

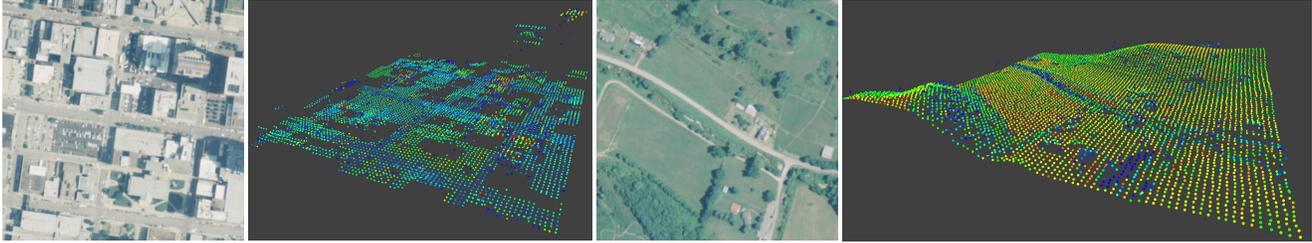


Figure 2: Examples of our multi-modal dataset. A geospatially aligned overhead image and corresponding point cloud are shown for an urban scene (left) and a rural scene (right). Point cloud coloring represents the relative grayscale intensity.

### 3. A Multi-Modal Dataset for Free-Flow Speed Estimation

We introduce a large-scale dataset for free-flow speed estimation that combines free-flow speed data, point clouds obtained from airborne LiDAR, and overhead imagery. Our dataset extends a recently introduced dataset that relates speed data on road segments throughout Kentucky, USA with overhead imagery. We begin by giving an overview of this existing dataset, then describe how we augment it with geospatially consistent 3D point cloud data.

#### 3.1. Kentucky Free-Flow Speed Dataset

The Kentucky Transportation Center [1] licensed and aggregated HERE Technologies’s speed data across uncongested periods to produce free-flow speeds for road segments across Kentucky. The speed data was then spatially joined with the Kentucky Transportation Cabinet’s highway inventory data. For each road segment, Song et al. [23] collected an overhead image centered at the location of the free-flow speed label. The overhead imagery is from the National Agriculture Imagery Program (NAIP) with 1m ground-sample distance (GSD). A single image has a spatial coverage of  $200 \times 200m^2$ . Each image was resized to  $224 \times 224$  pixels and rotated to ensure the road segment was aligned with direction of travel to the North. The dataset is representative of rural, urban, highway and arterial roads ranging in structure from multi-lane paved roads to single-lane dirt/gravel roads.

#### 3.2. Augmenting with Point Cloud Data

We augment this dataset with 3D point clouds extracted from LiDAR data collected by the Kentucky Division of Geographic Information’s KyFromAbove [9] program. Unlike overhead images, geometric features such as change in elevation, road curvature, lane delineation markings, lane width, proximity to neighboring structures, and more, can be easily detected from airborne LiDAR point clouds. The LiDAR data was stored as a collection of  $1524 \times 1524m^2$  tiles covering the state of Kentucky. To relate point cloud data with geospatially aligned overhead imagery and free-

flow speed data, we performed a two-step process consisting of LiDAR tile selection and point cloud sampling.

In order to associate each free-flow speed label with its containing tile, we constructed an R-tree using each tile’s geospatial coordinates. Then for each tile, we constructed a k-d tree over a random subset of points (50% selected uniformly at random) to support faster nearest-neighbor lookup. To generate an aligned point cloud, we use an  $80 \times 80$  uniformly sampled grid to guide point subsampling in a bounding box of the same spatial dimension as the overhead image. The resulting point cloud is centered on the target label location and is used to represent the spatial features of a given road segment.

An overview of our dataset is shown in Figure 2. The urban road segment depicted in Figure 2 (left) corresponds to the point cloud of the same road segment. The point cloud intensities are illustrated by dark blue roadways in stark contrast with the red roof tops of sky scrapers (top right). Similarly, the rural road segment point cloud in Figure 2 (right) shows the dynamic topography of the surrounding landscape not present in the corresponding overhead image.

## 4. Methods

We introduce *RasterNet*, an architecture for free-flow speed estimation that fuses multi-modal sensory input from overhead images and 3D point clouds. A visual overview of our architecture is given in Figure 3. Overhead images pass through an image encoder, while point clouds and raster center locations are passed through a point cloud encoder. A set of raster center locations guide point cloud feature extraction to produce geospatially consistent features between the two domains. The two sets of features are then channel-wise concatenated before being passed through a shared model to produce a free-flow speed prediction. We describe each component of our architecture in detail in the following sections.

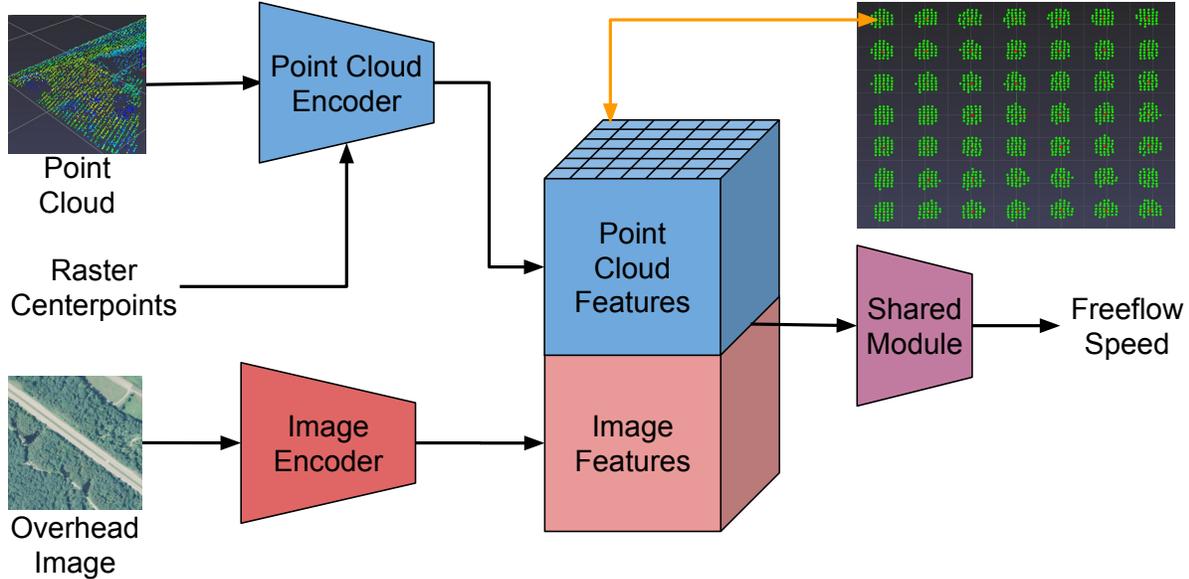


Figure 3: An overview of the *RasterNet* architecture. Overhead images pass through an image encoder, while point clouds and raster center locations are passed through a point cloud encoder. Each cell of the point cloud feature map corresponds to a set of features of a local point cloud neighborhood. The two sets of features are channel-wise concatenated before being passed through a shared model (ResNet block) to produce a free-flow speed prediction.

#### 4.1. Learning Visual Features

*RasterNet*'s image encoder is based on ResNet [6], a popular neural network architecture that contains residual connections. Specifically, we chose ResNet18 for our image feature extractor due to its low parameter count and relatively high performance on other tasks such as ImageNet [4] classification. In this work, we truncate before the average pooling layer such that the final encoding is size  $C \times H \times W$ , where  $C$  refers to the channel dimension and  $H$  and  $W$  refer to the spatial dimensions of the output feature map.

#### 4.2. Extracting Point Cloud Features

We explore two strategies for extracting point cloud features: (1) using a learning-based method (*RasterNet Learn*), and (2) using features computed from structural statistics (*RasterNet Statistics*). We begin by describing how we define a grid of point locations to align point cloud features with visual features.

##### 4.2.1 Aligning Visual and Point Cloud Features

An inherent challenge of training deep learning models on point clouds is their lack of fixed and consistent structure. To guide point cloud feature extraction, we propose a structural tool, the raster center grid, to impose consistent structure on extracted point cloud features. As Figure 4 illustrates, each raster center (red dots) binds local neighborhoods of point cloud features to a fixed location in a  $H \times W$

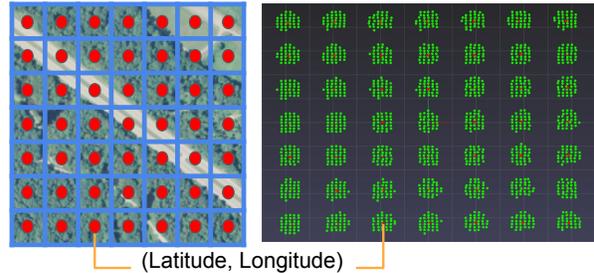


Figure 4: Image features are paired with point cloud features using a grid of raster center points (red dots), ensuring geospatial consistency between the two feature sets.

grid, similar to how CNNs group image features. The raster center grid was constructed to geospatially align with the pixel locations of the 2D image encoding. We did this by linearly sampling an  $H \times W$  grid within the known bounding box of the overhead image. This enables features extracted from point clouds to be directly paired with image features in a geospatially consistent manner.

##### 4.2.2 Learned Features

The *RasterNet Learn* model uses a modified PointNet++ [18] architecture as a learned point cloud feature extractor. PointNet++ was selected as a point cloud feature extractor because of its simplicity and high accuracy on point

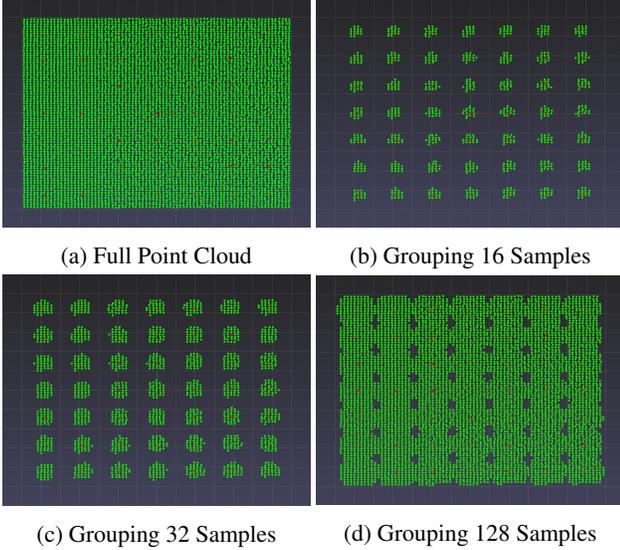


Figure 5: PointNet++ [18] style multi-scale grouping depicted for a point cloud (a) centered on a known free-flow speed label location. Grouping operations are performed around each of the raster centers (red) at different scales and number of samples. Local point clouds (green) are grouped at varying sample sizes: 16 samples (b), 32 samples (c), and 128 samples (d).

cloud tasks. The publicly available PyTorch [16] implementation of PointNet++ from Wijmans [26] was modified so the second multi-scale grouping layer performed grouping around the raster center grid of a given point cloud instead of using furthest point sampling. This modification allows the point cloud features to be combined with image features while maintaining spatial consistency. After the second multi-scale grouping layer the remainder of Pointnet++ was replaced with a series of  $1 \times 1$  convolutions that reduced the number of collected features per raster center to 16.

### 4.2.3 Statistical Features

Alternatively, we also developed the *RasterNet Statistics* model that directly extracts structural statistic features from the input point cloud. The *RasterNet Statistics* model replaced the PointNet++ architecture from *RasterNet Learn* model with a single instance of multi-scale grouping, as depicted in Figure 5. This approach allowed the model to aggregate spatial features at small, medium, and large scales. A single-scale grouping operation collects groups of points around each of the raster center points. Multi-scale grouping transformed each input point cloud into three separate collections point clouds, each for a different neighborhood group size  $k$ .

Inspired by Liu et al’s [12] work on place recognition us-

Table 1: Structural Statistics

Structural Statistic	Equation
<i>Change of Curvature</i>	$C = \frac{\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3}$
<i>Omni-variance</i>	$O = \frac{\sqrt[3]{\lambda_3}}{\lambda_1 + \lambda_2 + \lambda_3}$
<i>Linearity</i>	$L = \frac{\lambda_1 - \lambda_2}{\lambda_1}$
<i>Eigenentropy</i>	$A = - \sum_{j=1}^3 \lambda_j \ln \lambda_j$
<i>Local Point Density</i>	$D = \frac{k}{\frac{4}{3} \prod_{j=1}^3 \lambda_j}$
<i>2D Scattering</i>	$S_{2D} = \lambda_1^{2D} + \lambda_2^{2D}$
<i>2D Linearity</i>	$L_{2D} = \frac{\lambda_2^{2D}}{\lambda_1^{2D}}$
<i>Verticality</i>	$V = v_{3,z}$
<i>Max Height Difference</i>	$\Delta Z = \max(x_z) - \min(x_z)$
<i>Height Variance</i>	$\sigma^2 = \frac{1}{N} \sum_{i=1}^N (x_{z_i} - \bar{x}_z)^2$

ing LiDAR point cloud structural features, we extracted statistical features from airborne LiDAR point clouds. Let  $x_{i,k}$  be a point cloud containing  $k$  neighborhood points around point  $i$ . Neighborhood statistical features were extracted by first calculating the covariance matrix of  $x_{i,k}$ . We compute eigenvalues of the covariance matrix, resulting in three eigenvalues  $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ . The structural statistics of the point cloud were calculated according to the equations listed in Table 1. Note, 2D statistics (Scattering and Linearity) were calculated using 2D eigenvalues, which were calculated from the covariance matrix of the 3D point cloud projected to the  $xy$ -plane. We use  $x_z$  to specify that we consider only the  $z$  component of the points in point cloud  $x$ , and  $x_{z_i}$  to express the  $i$ th point in  $x_z$ . For verticality [25],  $v_{3,z}$  is the  $z$  component of the eigenvector corresponding to the smallest eigenvalue,  $\lambda_3$ .

In order to get statistical features at local point cloud regions, we extracted 10 statistical features for each of the local point cloud neighborhoods corresponding to each raster center. We compute this for three different group sizes  $k \in \{16, 32, 128\}$ , resulting in 30 total features per local point cloud neighborhood. The structure features of each raster center are tiled to create a single  $30 \times H \times W$  feature map, corresponding to the feature map resolution of the image encoding. This is then reduced to  $10 \times H \times W$  by applying a series of  $1 \times 1$  convolutions.

### 4.3. Feature Fusion for Estimating Free-Flow Speed

The visual features and point cloud features are channel-wise concatenated and passed through a shared module whose role is to extract high-level features from the combined domains and produce a free-flow speed estimate. The spatial correspondence established by the raster center grid between the overhead image and point cloud features ensured that the two sets of input features are spatially aligned. To represent the shared module, we use a single ResNet18 block and a drop out layer for regularization followed by a fully connected layer with  $K$  outputs.

### 4.4. Implementation Details

We model the free-flow speed prediction as a multi-class classification problem. Free-flow speeds were binned into  $K = 79$  possible classes, each in 1mph increments. Our models train using the cross-entropy loss ( $L$ ) with a *softmax* activation defined as follows,

$$L(Y, \hat{Y}) = -\frac{1}{N} \sum_{i=1}^N \log \left( \frac{e^{y_i}}{\sum_j^K e^{\hat{y}_{i,j}}} \right). \quad (1)$$

Let  $y_i \in Y$  be a positive class bin label for the  $i$ th sample from  $N$  training samples. The predicted probability from the distribution  $\hat{Y}$  for the  $i$ th sample from the  $j$ th class was expressed as  $\hat{y}_{i,j}$ , where  $j \in \{0, 1, \dots, K\}$ .

The x and y dimensions of point clouds in dataset were translated such that the origin corresponds to the center of the matching overhead image. The height dimension of the point cloud was normalized by subtracting from the median height for the given point cloud and the point intensity values were normalized by dividing by 255. All point clouds were then rotated such that the direction of travel of the target road was pointed north, similar to how the imagery was aligned.

Given an input image of size  $224 \times 224$ , the output feature map of the image encoder is of size  $C \times 7 \times 7$ . Therefore, we define the raster center grid to be of size  $H \times W = 7 \times 7$ . Unlike the shared module and point cloud encoder, the image encoder was pretrained on ImageNet [4] and frozen. The training configuration of each network included an Adam optimizer with learning rate of  $1 \times 10^{-6}$  and weight decay of 0.1. The learning rate was reduced by a factor of 10 every 25 epochs.

## 5. Evaluation

We present an ablation study, a quantitative analysis compared with an existing approach on a held-out test set, and a qualitative evaluation of our best method compared with known free-flow speeds. Training, validation, and test dataset partitioning followed the methodology established by Song et al. [23]. Each model was evaluated on the set of

Table 2: Free-flow Speed Estimation Model Performances

Method	Accuracy
<i>Song et al. Image Only [23]</i>	37.60%
<i>Song et al. Image + Road Features [23]</i>	49.86%
<i>Reduced PointNet++ [18]</i>	34.08%
<i>ResNet [6]</i>	42.01%
<i>RasterNet Statistics</i>	47.75%
<b><i>RasterNet Learn</i></b>	<b>50.47%</b>

weights chosen based on the lowest validation loss. Roads within the borders of the following Kentucky, USA counties were held-out for the test set: Bell, Lee, Ohio, Union, Woodford, Owen, Fayette, and Campbell. The validation set was constructed from 1% of the training set samples.

### 5.1. Quantitative Evaluation

We performed an ablation study comparing different image feature extractors and the impact of point cloud features on free-flow speed estimation, shown in Table 2. Following previous work, free-flow speed estimation was evaluated using within-5mph accuracy. In this metric, predicted free-flow speed is considered correct if it is within 5mph of the true speed.

We evaluated the performance of a full ResNet model trained only on overhead imagery in order to highlight the differences in image feature extractors compared to previous work. Specifically, we compared an Xception-style [23] architecture to our ResNet18 architecture. The first 3 blocks and the 4th block’s residual sub-block were frozen, similar to the *RasterNet* architectures. The smaller ResNet (12M parameters) network trained only with image features outperformed the Xception-based network (23M parameters) by 5% average within-5mph test accuracy, suggesting it was the superior image feature extractor for this task.

To understand the impact of augmenting point cloud features with visual features, we compared our approach to a point cloud only baseline that uses a reduced PointNet++ model as in *RasterNet Learn*. Following the same strategy, the second multi-scale grouping layer was modified to extract features at raster center locations. The number of fully connected layers in the last MLP (after the multi-scale grouping layer) was reduced to two layers for faster training. The reduced PointNet++ with raster center locations had the worst performance of all of the evaluated models. While the performance is still respectable, it shows that point clouds alone do not provide the features necessary for this task.

Next, we examined the performance impact of the point cloud feature extraction strategies. We observed that the learned features (*RasterNet Learn*) perform better than the structural features (*RasterNet Statistics*). By combining fea-

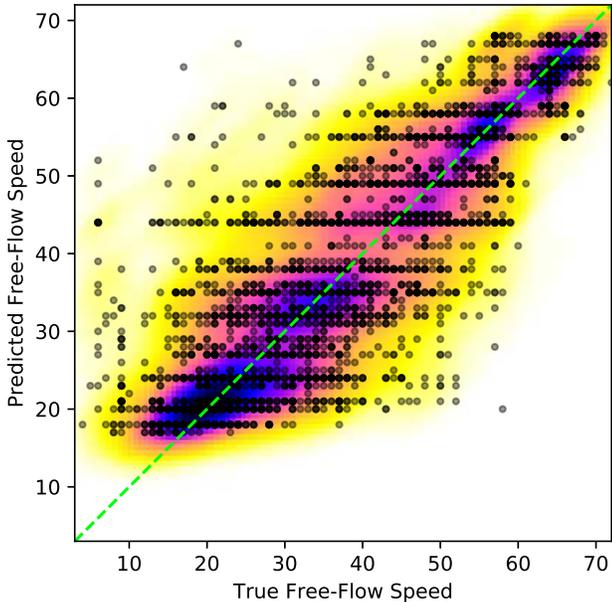


Figure 6: Scatterplot of free-flow speed predictions on the test set from the *RasterNet Learn* model compared with known speed labels. Overlaid heatmap depicts higher point density in darker color. Optimal performance should follow the green line.

tures from both point cloud and overhead imagery, we are able to greatly improve the accuracy compared to the single modality networks. Furthermore, our *RasterNet Learn* model achieves state-of-the-art performance over the previous best method, despite not using highway geometric features. In subsequent experiments, we use the *RasterNet Learn* model.

## 5.2. Qualitative Evaluation

Figure 6 shows a scatter plot of the model’s prediction versus ground-truth free-flow speeds on the test set. In addition, it includes a heatmap, generated using kernel density estimation, to make the joint distribution clear. Overall, the highest density of predictions (the darker colors) follow the green line, indicating a positive relationship with the true free-flow speeds. While the model had difficulties in predicting speeds accurately for roads with true free-flow speeds  $< 10$ mph, for most other roads the model predicts speeds close to the ground-truth speed.

Additionally, we visualized the *RasterNet Learn* model by constructing free-flow speed maps. We generated these maps with the ground truth and predicted free-flow speeds for 3 Kentucky counties from the test set: Fayette, Woodford, and Union. Since Fayette and Woodford counties are adjacent, we visualize them on the same map in Figure 7 (a) and (b). Figure 7 (b) suggests that the model is capable of estimating free-flow speeds on highways accurately, as

shown by two major highways both being red in both maps. Unlike highways and suburban areas, urban arterial road segments, as seen in the Lexington city center of Figure 7 (a) and (b), are more challenging. These low speed urban arterial road segments are impacted by traffic signal timings which play a large role in regulating vehicle speeds, which are not captured in overhead imagery and LiDAR data.

The model performs well in rural counties, such as Union county in Figure 7 (c) and (d), with speeds primarily ranging from 30-50mph. Note in Figure 7 (c), the road segment on the far left is dark blue, indicating free-flow speeds  $< 20$ mph. The predicted free-flow speed map in Figure 7 (d) suggests that the model predicts speeds  $> 20$ mph for said road segment. The road segment in question is a dirt road, an underrepresented road type in the training set, likely causing the poor performance in this scenario.

## 6. Conclusion

We presented a novel multi-modal architecture for free-flow speed estimation, *RasterNet*, that jointly processes aligned overhead images and corresponding 3D point clouds from airborne LiDAR. To support training and evaluating our methods, we introduced a large dataset of free-flow speeds, overhead imagery, and LiDAR point clouds across the state of Kentucky. We evaluated our approach on a benchmark dataset, achieving state-of-the-art results without requiring explicit highway geometric features, unlike the previous best method. Additionally, we show how our approach can be used to generate large-scale free-flow speed maps, a potentially useful tool for transportation engineering and roadway planning. Our results demonstrate that a combination of overhead imagery and 3D point clouds can replace and ultimately outperform existing approaches that rely on manually annotated input data. Our hope is that our dataset and proposed approach will inspire future work in estimating free-flow speeds from multi-modal input data.

## References

- [1] Mei Chen, Xu Zhang, and Eric R Green. Analysis of historical travel time data. Technical report, Kentucky Transportation Cabinet, 2015.
- [2] S Daneshmand, H Rastiveis, and B Hosseiny. Cnn-based feature-level fusion of very high resolution aerial imagery and lidar data. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences*, 2019.
- [3] Matthew D Deardoff, Brady N Wiesner, and Joseph Fazio. Estimating free-flow speed from posted speed limit signs. *Procedia-social and behavioral sciences*, 16, 2011.
- [4] Jia Deng, Wei Dong, R. Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [5] Renaud Dubé, Daniel Dugas, Elena Stumm, Juan Nieto, Roland Siegwart, and Cesar Cadena. Segmatch: Segment

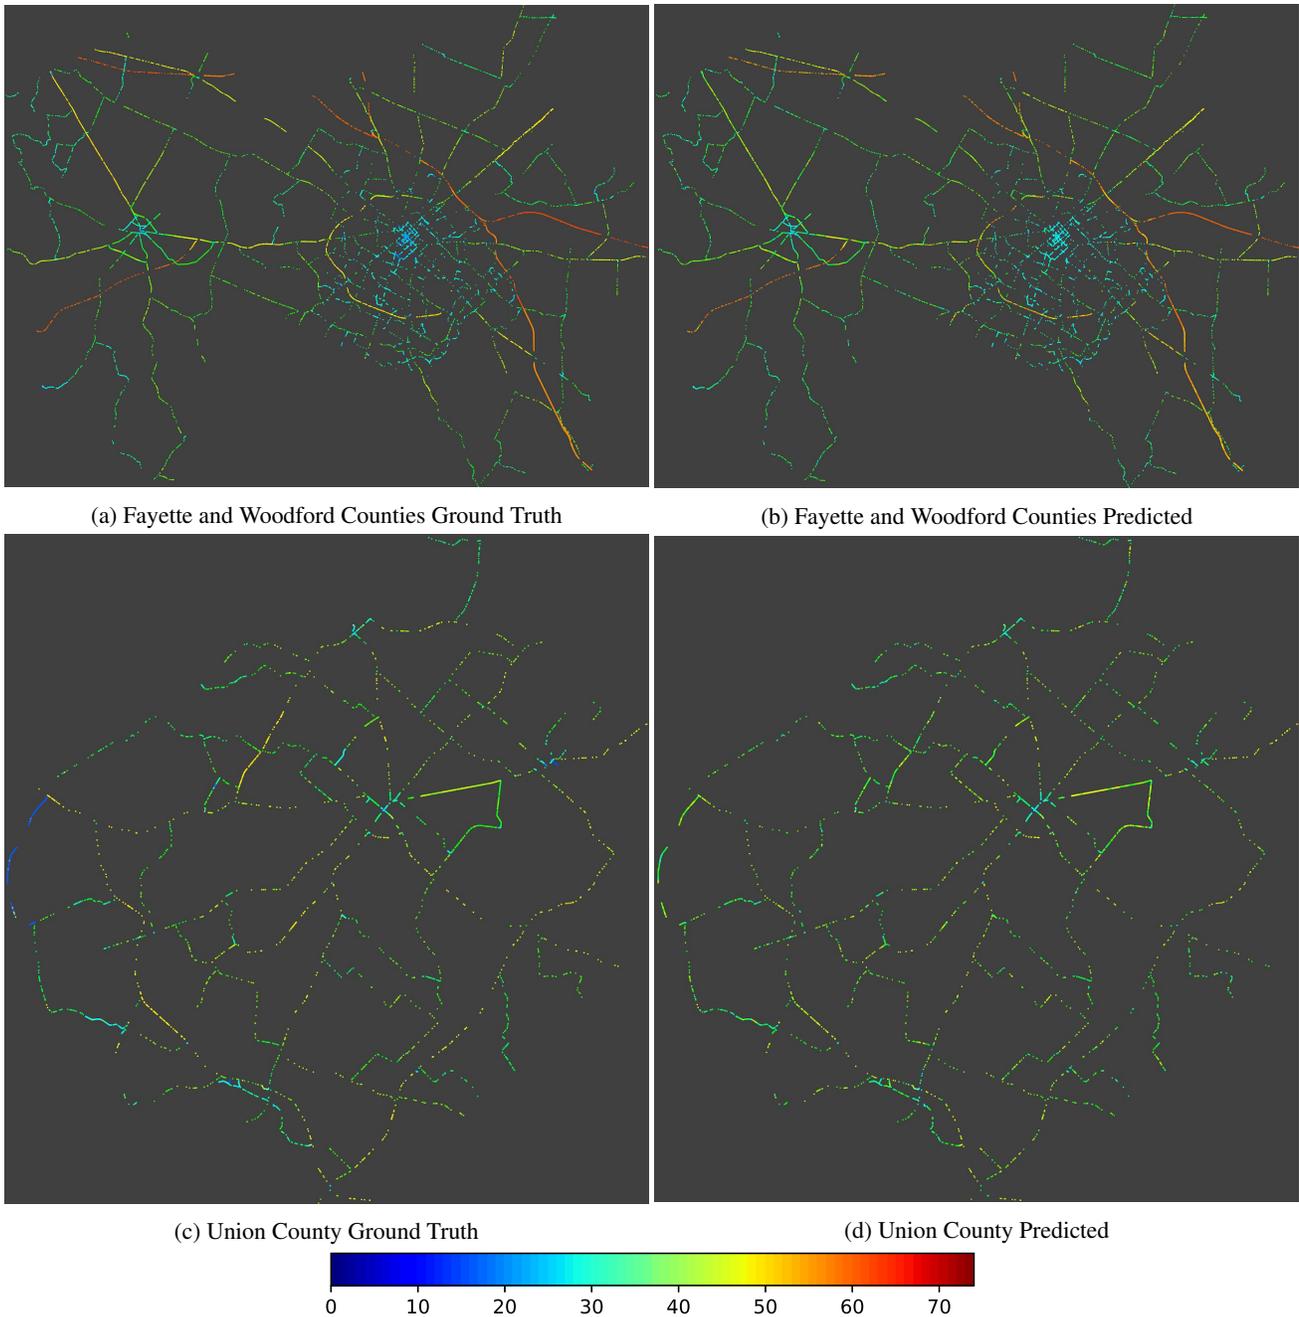


Figure 7: Ground truth and predicted speed maps for both Woodford (top left small city), Fayette (top larger city) and Union (bottom) counties in Kentucky, USA.

based place recognition in 3d point clouds. In *International Conference on Robotics and Automation*, 2017.

[6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

[7] Tingting Huang. Traffic speed estimation from surveillance video data: For the 2nd nvidia ai city challenge track 1. In

*IEEE International Conference on Computer Vision Workshops*, 2018.

[8] Maximilian Jaritz, Tuan-Hung Vu, Raoul de Charette, Émilie Wirbel, and Patrick Pérez. xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation, 2019.

[9] Kentucky Division of Geographic Information

- KyFromAbove. <http://kyfromabove-kygeonet.opendata.arcgis.com/>.
- [10] Shiyi Lan, Ruichi Yu, Gang Yu, and Larry S Davis. Modeling local geometric structure of 3d point clouds using geocnn. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [11] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *European Conference on Computer Vision*, 2018.
- [12] Zhe Liu, Shunbo Zhou, Chuanzhe Suo, Yingtian Liu, Peng Yin, Hesheng Wang, and Yunhui Liu. Lpd-net: 3d point cloud learning for large-scale place recognition and environment analysis. In *International Conference on Computer Vision*, 2019.
- [13] Highway Capacity Manual. Hcm2010. *Transportation Research Board, National Research Council, Washington, DC*, page 1207, 2010.
- [14] Jiageng Mao, Xiaogang Wang, and Hongsheng Li. Interpolated convolutional networks for 3d point cloud understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [15] Jiageng Mao, Xiaogang Wang, and Hongsheng Li. Interpolated convolutional networks for 3d point cloud understanding. In *International Conference on Computer Vision*, 2019.
- [16] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.
- [17] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [18] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, 2017.
- [19] Nagui M Roupail, SangKey Kim, and Seyedbehzad Aghdashi. Application of high-resolution vehicle data for free-flow speed estimation. *Transportation Research Record*, 2615(1):105–112, 2017.
- [20] Ch Ravi Sekhar, J Nataraju, S Velmurugan, Pradeep Kumar, and K Sitaramanjaneyulu. Free flow speed analysis of two lane inter urban highways. *Transportation research proceedings*, 17:664–673, 2016.
- [21] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *CVPR*, 2019.
- [22] Ary P Silvano and Karl L Bang. Impact of speed limits and road characteristics on free-flow speed in urban areas. *Journal of transportation engineering*, 142(2):04015039, 2016.
- [23] Weilian Song, Tawfiq Salem, Hunter Blanton, and Nathan Jacobs. Remote estimation of free-flow speeds. In *IEEE International Geoscience and Remote Sensing Symposium*, 2019.
- [24] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. Kpconv: Flexible and deformable convolution for point clouds. *International Conference on Computer Vision*, 2019.
- [25] Martin Weinmann, Boris Jutzi, and Clément Mallet. Semantic 3d scene interpretation: A framework combining optimal neighborhood size selection with relevant features. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2(3):181, 2014.
- [26] Erik Wijmans. Pointnet++ pytorch. [https://github.com/erikwijmans/Pointnet2\\_PyTorch](https://github.com/erikwijmans/Pointnet2_PyTorch), 2018.