

The following publication Z. Xu, Y. Sun and M. Liu, "Topo-Boundary: A Benchmark Dataset on Topological Road-Boundary Detection Using Aerial Images for Autonomous Driving," in IEEE Robotics and Automation Letters, vol. 6, no. 4, pp. 7248-7255, Oct. 2021 is available at <https://doi.org/10.1109/LRA.2021.3097512>.

Topo-boundary: A Benchmark Dataset on Topological Road-boundary Detection Using Aerial Images for Autonomous Driving

Zhenhua Xu, *Student Member, IEEE*, Yuxiang Sun, *Member, IEEE*, and Ming Liu, *Senior Member, IEEE*

Abstract—Road-boundary detection is important for autonomous driving. It can be used to constrain autonomous vehicles running on road areas to ensure driving safety. Compared with online road-boundary detection using on-vehicle cameras/Lidars, offline detection using aerial images could alleviate the severe occlusion issue. Moreover, the offline detection results can be directly employed to annotate high-definition (HD) maps. In recent years, deep-learning technologies have been used in offline detection. But there still lacks a publicly available dataset for this task, which hinders the research progress in this area. So in this paper, we propose a new benchmark dataset, named *Topo-boundary*, for offline topological road-boundary detection. The dataset contains 25,295 1000×1000 -sized 4-channel aerial images. Each image is provided with 8 training labels for different sub-tasks. We also design a new entropy-based metric for connectivity evaluation, which could better handle noises or outliers. We implement and evaluate 3 segmentation-based baselines and 5 graph-based baselines using the dataset. We also propose a new imitation-learning-based baseline which is enhanced from our previous work. The superiority of our enhancement is demonstrated from the comparison. The dataset and our-implemented code for the baselines are available at <https://tonyxuqaq.github.io/Topo-boundary/>.

Index Terms—Road-boundary Detection, Imitation Learning, Large-scale Dataset, Autonomous Driving.

I. INTRODUCTION

ROAD boundary refers to the dividing line between the road area and off-road area. It can be used to constrain self-driving cars running on road areas, which is important to the safety of autonomous driving. Currently, most existing works rely on vehicle-mounted sensors (e.g., cameras or Lidars) [1]–[4] for online road-boundary detection. However, online detection could be severely degraded by occlusions, which is very common in real road environments. Moreover, online detection could be restricted by limited computing resources on vehicles, especially for deep-learning models that require GPU. To relieve the aforementioned issues, some works resort

to detecting road boundary (or its analogies, such as lane or curb) offline using bird-eye-view (BEV) point-cloud maps or aerial images. As the results are presented in BEV images, they can also be directly utilized for annotating high-definition (HD) maps. So we detect road boundaries offline in this paper. Moreover, our detection results are presented in the form of topological graphs (i.e., vertices and edges). This is important because besides pixel-level annotations, the topological graphs can also identify road-boundary instances and find the spatial connection information of the boundaries.

The published literature working on topological road-boundary detection is very limited. Most of existing works focus on the analogy problem: line-shaped object detection, such as road-lane, road-curb and road-network detection [5]–[11]. They can be generally divided into two categories: segmentation-based solutions and graph-based solutions. A typical pipeline of the former is first using semantic segmentation to get rough detection results, and then performing heuristic post-processing algorithms on the segmentation maps to refine the results. The latter directly finds the graph structure for the line-shaped objects through iterative graph growing. With the great advancement of artificial intelligence, deep-learning technologies are adopted by current existing methods, and great superiority over traditional algorithms is achieved. However, deep-learning-based methods require large-scale datasets to train a model. To the best of our knowledge, there is still no publicly available datasets for road-boundary detection in BEV images, which hinders the research process in this area. This motivates us to build a large-scale benchmark dataset, named *Topo-boundary*, for road-boundary detection.

Our dataset is derived from the *NYC Planimetric Database* [12]. The raw geographic data (a large whole map, see Fig. 1) in the *NYC Planimetric Database* covers the whole New York City, including 2,147 5000×5000 -sized 4-channel (i.e., red, green, blue and an infrared channel) aerial image tiles (i.e., small squared aerial image constituting the large whole map) and a set of polylines as the road-boundary ground-truth label. To build our *Topo-boundary*, we convert the polyline label to graph (i.e., vertices and edges) and split every tile into 25 smaller 1000×1000 -sized patches. After removing patches with inappropriate annotations, such as patches without road boundaries, 25,295 patches remain. Each patch consists of a 4-channel image and 8 labels for different sub-tasks, such as binary semantic segmentation, orientation learning [6], etc. The dataset can be used for both segmentation-based solutions and graph-based solutions for road-boundary detection.

To facilitate future research on road-boundary detection, we implement and evaluate multiple baselines, including 3 segmentation-based baselines, 5 graph-based baselines. We also

Manuscript received February 24, 2021; Revised May 23, 2021; Accepted Jun 21, 2021. This paper was recommended for publication by Editor Pauline Pounds upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by Collaborative Research Fund by Research Grants Council Hong Kong, under Project No. C4063-18G, Department of Science and Technology of Guangdong Province Fund, under Project No. GDST20EG54 and Zhongshan Municipal Science and Technology Bureau Fund, under project ZSST21EG06, awarded to Prof. Ming Liu. (*Corresponding author: Ming Liu.*)

Zhenhua Xu is with the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong (email: zxubg@connect.ust.hk).

Yuxiang Sun is with the Department of Mechanical Engineering, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong (email: yx.sun@polyu.edu.hk, sun.yuxiang@outlook.com).

Ming Liu is with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong (email: eelium@ust.hk).

Digital Object Identifier (DOI): see top of this page.

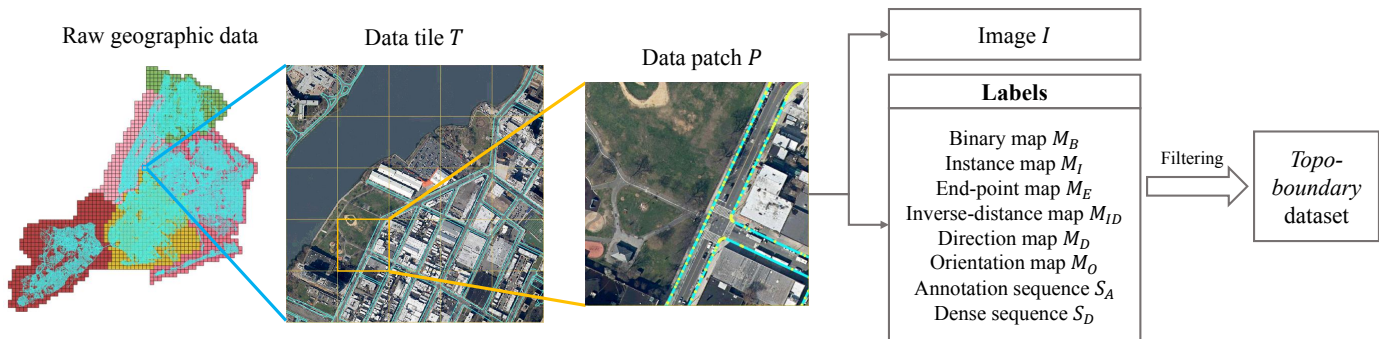


Fig. 1: The pipeline to create our dataset from the *NYC Planimetric Database*. The raw geographic data contains 2,147 data tiles. The different colors represent the 5 boroughs and water areas in the New York City. Each tile (T) is with 5000×5000 size. We split each tile into 25 smaller patches ($T = \{P_i\}_{i=1}^{25}$). Each patch (P) contains a 4-channel 1000×1000 -sized aerial image I and ground-truth road boundaries G_{gt} . The G_{gt} in the figure is annotated by polylines, of which edges and vertices are denoted by cyan lines and yellow points, respectively. Based on G_{gt} , we generate 8 types of ground-truth labels for different tasks. The filtering step is to remove patches based on pre-defined rules. After filtering, 25,295 patches finally remain in our *Topo-boundary*. For better visualization, only RGB channels of the patch image are visualized. Moreover, the width of the ground-truth polylines is increased (the actual width is one pixel). This figure is best viewed in color.

design and compare a new imitation-learning-based baseline which is enhanced from our previous work *iCurb* [11]. For all the baselines, the input is a 4-channel aerial image, the output is the graph representing road boundaries. Our implemented code for these baselines are open-sourced along with our dataset.

As line-shaped objects are usually long, thin and irregular, only using pixel-level metrics is not sufficient to evaluate the performance of a method. In past works, average path length similarity (APLS) has been widely used to evaluate topology correctness [13]. But since in most cases road boundaries are simple polylines without branches, APLS suffers from randomness and inefficiency. Inspired by the naive connectivity metric proposed in [14], we design a new entropy-based connectivity metric (ECM) to evaluate the connectivity of the obtained road-boundary graph, which is effective and more efficient. We summarize our contributions here:

- 1) We release the *Topo-boundary* benchmark dataset for topological road-boundary detection using aerial images. To the best of our knowledge, this is the first publicly available dataset for this task.
- 2) We propose new evaluation metrics for the task, including relaxed pixel-level metrics and a new entropy-based connectivity metric.
- 3) We quantitatively evaluate 9 baseline models using the proposed dataset, which could facilitate comparative study for future methods.
- 4) We open-source our-implemented code for the baselines. The code can also be modified with a little effort for other line-shaped object detection.

II. RELATED WORKS

A. Segmentation-based line-shaped object detection

There are very few segmentation-based works directly detecting road boundaries topologically. So we review several line-shaped object detection works [5], [6], [15]. Volodymyr *et al.* [15] proposed the first deep-learning model with iterative refinement for road-network detection using aerial images. Their

solution was further improved by Anil *et al.* [6], in which the authors not only proposed better networks and a training scheme, but also utilized the orientation map to enhance the semantic segmentation results. Different from the above papers using iterative refinement, the solution proposed by Mattyus *et al.* [5] first put forward connection candidates to correct disconnections, and then trained another network to filter candidates. Although segmentation-based solutions are efficient to carry out, they can only produce results at the pixel-level. Moreover, they were often implemented with multi-stage pipelines, so they could not be optimized as a whole, making the graph sensitive to incorrect topology.

B. Graph-based line-shaped object detection

Taken images as input, graph-based methods can directly output the graph representing target objects. Iterative graph growing is the most commonly used technique, which generates vertices along with the line-shaped object starting from an initial vertex. Past works focusing on other line-shaped objects are similar to our task, such as road-lane detection [9], [16], road-network detection [7], [8], [17] and road-curb detection [11]. *RoadTracer* [7] is believed to be the first work in this category. In [7], the authors trained a multi-layer CNN and successfully extracted the graph of very large road networks. Liang *et al.* [14] directly worked on road-boundary detection in the BEV point-cloud map. They facilitated semantic segmentation by direction map prediction and proposed cSnake for iterative graph growing. To better solve the graph growing task, line-shaped object detection was first analyzed from the perspective of imitation learning in our previous work *iCurb* [11]. *iCurb* presented superiority over other works on road curb detection owing to the training strategy and graph growing policy.

C. Evaluation metrics for topological correctness

In the past works on road network detection, topological correctness was usually measured by shortest-path-based metrics, such as APLS [13] and too long/too short (TLTS) similarity [18].




Fig. 2: Visualization of label maps for a patch. (a) 4-channel patch image I (only RGB is visualized here); (b) Binary map M_B ; (c) Instance map M_I . Different road boundary instances are labeled with different gray values; (d) End-point map M_E ; (e) Inverse-distance map M_{ID} . Each pixel value is the reciprocal of its shortest distance to the boundary; (f) Direction map M_D . The red channel is the Sobel derivative of columns, the green channel is the Sobel derivative of rows, and the blue channel is the sum of the other two channels; (g) Orientation map M_O . Pink arrows are the schematic demonstration of directional information of the road boundary. For better visualization, the line width in (b) and (c) is increased (the actual width is one pixel). This figure is best viewed in color. Please zoom in for details. For more examples, please see our supplementary document.

However, due to the random sampling process, these metrics suffer from inefficiency and randomness. Alternatively, in [14], a naive metric measuring the connectivity of road boundaries was proposed, but it is too sensitive to noises and may cause incorrect evaluation results. *iCurb* modified the naive metric and addressed the problem to some extent. However, incorrect evaluation still happens. So in this work, we make further revisions to this metric and propose an entropy-based metric for connectivity evaluation.

D. Datasets for line-shaped object detection

To the best of our knowledge, [14] is the only published work on road-boundary detection using BEV images. In [14], a dataset was created, which contained BEV point-cloud images, RGB images and elevation gradient images. But unfortunately, the dataset was not publicly available. In the road-network detection area, the data was obtained from commercial geospatial datasets [13] or by processing raw geographic images collected from public platforms [19]. Although these datasets are publicly available, due to the topological differences between road networks and road boundaries, they could not be used in our task.

III. THE PROPOSED DATASET

To create our *Topo-boundary* dataset, we generate images and ground-truth labels from the raw geographic data provided by the *NYC Planimetric Database* [12]. The WGS84 coordinate system is used in [12], while in our *Topo-boundary*, all the coordinates are transformed to the image coordinate system to facilitate research in this area. Fig. 1 shows the processing pipeline to create our *Topo-boundary*. We first split large data tiles into patches, then generate ground-truth labels for each patch, and finally remove inappropriate patches according to predefined rules.

A. Splitting data tiles into patches

The raw geographic data covers the whole area of the New York City, including 2,147 data tiles. Each tile (T) is a 5000×5000 -sized image with road-boundary ground-truth label. Each pixel represents 1 foot (around 15.2cm). We split each tile into 25 smaller patches ($T = \{P_i\}_{i=1}^{25}$) to reduce the GPU memory cost during training. Since our task is mainly performed at the

patch level, the subscript i is removed from P_i for expression conciseness in the following text. Each patch (P) contains a 4-channel (red, green, blue, infrared, the infrared channel is more sensitive to vegetation and soil) 1000×1000 -sized aerial image I and ground-truth road boundaries G_{gt} . We use G_{gt} to generate 8 labels for different sub-tasks.

The raw ground-truth label for road boundaries is the `PAVEMENT_EDGE` feature from the *NYC Planimetric Database*. `PAVEMENT_EDGE` is manually annotated in the form of polylines (a kind of graph without branches), of which the vertices and edges are recorded using the WGS84 coordinate system. The polylines cover road-surface edges, airport runways and alleys. From the whole map `PAVEMENT_EDGE`, we obtain the ground-truth road boundary G_{gt} for each patch P .

Unlike tile images, `PAVEMENT_EDGE` cannot be directly cut and split because: (1) some edges $e = (v_a, v_b)$ may have one vertex v_a in a patch while the other vertex v_b in another patch; (2) some road boundaries may be cut into multiple pieces by different patches. For the former, we replace the vertex outside the patch (i.e., v_b) with the intersection point of e and the patch edge. For the latter, we split the corresponding road boundaries into multiple instances and update `PAVEMENT_EDGE`.

After processing `PAVEMENT_EDGE`, for each patch, the ground-truth road boundary is still a set of polylines $G_{gt} = \{G_{gt}^i\}_{i=1}^N$. Each polyline represents a road-boundary graph instance $G_{gt}^i = (V_i, E_i)$, where V_i and E_i are the set of vertices and edges in G_{gt}^i , respectively. Vertices V_i are recorded in a list, and then edges in a graph can be easily obtained by connecting two neighboring vertices in the list, thus the edge set E_i is omitted. Note that G_{gt}^i is sparsely distributed (i.e., adjacent vertices are not eight-neighboring to each other in the image coordinate system).

B. Generating ground-truth labels for each patch

For each patch, we generate 8 different types of labels for different tasks. The labels are: annotation sequence S_A , dense sequence S_D , binary map M_B , instance map M_I , endpoint map M_E , inverse-distance map M_{ID} , direction map M_D and orientation map M_O . All the label maps are generated from ground-truth road boundaries (G_{gt}) within the current patch by our processing algorithms, and are recorded in the image coordinate system of the current patch.




Fig. 3: The sequence densification process for a part of a sample patch. The orange pixels represent vertices in the annotation sequence S_A . They are interpolated to realize every two adjacent vertices eight-neighboring to each other. The interpolation result is the dense sequence S_D . S_A can be regarded as a subset of S_D , which contains only the key vertices. The figure is best viewed in color.

1) *Annotation sequence and dense sequence*: As the ground-truth road boundaries are annotated vertex-by-vertex as chain sequences, the vertices of a road boundary instance G_{gt}^i are ordered starting from an initial vertex to the end vertex. We call the ordered vertices of G_{gt}^i as the annotation sequence S_A . S_A is critical to generate other labels and can be applied to train graph-based solutions. However, graph-based solutions trained by S_A suffer from the *teacher forcing* problem [20], which is caused by the data distribution mismatch between the training and inference period. It is also analyzed in [7]. To address this problem, the training label for iterative graph generation should be generated on-the-fly. Therefore, S_A needs to be densified. For every two adjacent vertices in S_A , interpolation is conducted to obtain the dense sequence S_D . In S_D , any two adjacent vertices are eight-neighboring to each other. The densification process is illustrated in Fig. 3.

2) *Binary map, instance map and endpoint map*: Binary map M_B is commonly used for semantic segmentation. The dense sequence S_D can be directly used to generate M_B by setting all the pixels covered by S_D as 1 (i.e., foreground pixels) and the other pixels as 0 (i.e., background pixels). In this way, we can get the binary map of a patch and its topological correctness can be guaranteed. The instance map M_I can be obtained similarly, but pixels covered by different instances are multiplied with instance IDs.

To facilitate segmentation and generate candidate initial vertices for iterative graph generation, endpoint prediction is required in some solutions [11], [14]. For each road boundary instance G_{gt}^i in a patch, there are 2 endpoints. We multiply each endpoint by a Gaussian kernel function.

3) *Inverse-distance map and direction map*: Inverse-distance map (M_{ID}) prediction is first proposed in [21] to facilitate semantic segmentation. The value of each pixel in this map is the reciprocal of its shortest distance to the ground-truth road boundary. Compared with the binary map M_B , M_{ID} is more informative considering that all the pixels of M_{ID} are encoded with information. The generation of M_{ID} is accelerated by GPU parallelism in our implementation.

To further make use of the spatial information in M_{ID} , the authors in [14] generate the direction map M_D based on M_{ID} . The direction map is a vector field ($M_D \in \mathbb{R}^{2 \times H \times W}$) of normal directions to the road boundary. It is calculated by

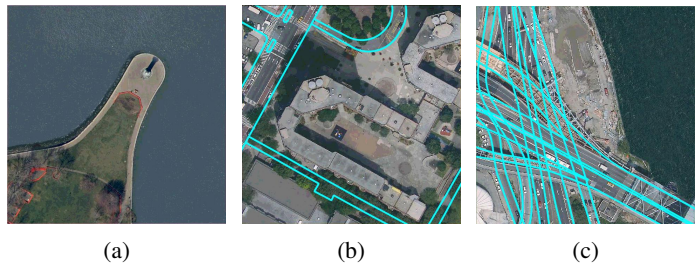


Fig. 4: Sample inappropriate patches. The cyan lines represent the ground-truth polylines for road boundaries. (a) A patch without road boundaries inside. Such case happens frequently near ocean and suburb areas; (b) A patch in which ground-truth road boundaries have intersection points. This is usually caused by alleys that snapped to the road boundary. Intersection points are not reasonable considering the topological characteristics of the road boundary; (c) A patch with very complicated scenarios. For these patches, few methods can obtain reasonable results, especially for graph-based methods. Therefore they are removed at this stage. For better visualization, only RGB channels of the patch image are visualized, and the width of the ground-truth polylines is increased (the actual width is one pixel).

taking the Sobel derivative of M_{ID} followed by a normalization step. Intuitively, M_D records the direction to the closest road boundary of each pixel by unit vectors. The application of M_D and M_{ID} encourages the neural network to focus on road boundaries.

4) *Orientation map*: Orientation learning [6] makes use of the direction and connection information of road networks, which greatly enhances the performance of semantic segmentation. Thus we provide the orientation map for this task to assist road-boundary detection. Due to the bi-direction nature of road boundaries, we regard them as undirected graphs. For each road boundary instance G_{gt}^i , we randomly select one end vertex as the starting point while the other as the ending point. Then from the starting point, for every two adjacent vertices, we calculate a directional vector and convert it to a radian value r . For all pixels covered by the edge connecting these two adjacent vertices, their values are set to r . Intuitively, the value of each foreground pixel is the radian value of the edge it belongs to. Fig. 2 displays all the label maps for a sample patch.

C. Removing patches according to pre-defined rules

We remove the inappropriate patches according to the following predefined rules:

- 1) The patch with no road boundary inside. Such a case is quite common, especially in suburb area images;
- 2) The patch with intersection points. In the raw ground-truth label of the road boundary, some alleys are snapped to it and cause intersections, which is not appropriate. Thus we remove these patches;
- 3) The patch that has very complex scenarios, such as overlapping interchanges.

The cases in the removed patches would happen in city-scale real-world applications. So, this could be treated as a limitation of our work at this stage. Fig. 4 shows sample inappropriate patches.

D. Data splitting

After filtering, there are finally 25,295 patches remaining in our *Topo-boundary*. We randomly split these patches into a training set (10,236 patches), a validation set (1,770 patches), a testing set (3,289 patches) and a pretraining set (10,000 patches). The pretraining set is used for multi-stage methods, like [5], or to pretrain the feature extraction module of the graph-based solutions to accelerate the training convergence, such as [9]. If pretraining is not needed, the samples in this set could be used for training. The boroughs of New York City (e.g., Queens or Manhattan) is not considered during data splitting. The above data splitting scheme is adopted in the experiments of this paper. Users are free to split the dataset according to their needs.

IV. EVALUATION METRICS

To ensure comprehensive and fair comparison, we design our evaluation metrics, including 3 relaxed pixel-level metrics (i.e., Precision, Recall and F1-score), the naive connectivity metric [14], APLS [13] and an entropy-based connectivity metric (ECM) for topological correctness measurement. For all the metrics, larger values indicate better performance.

A. Pixel-level metrics

Pixel-level metrics (i.e., Precision, Recall and F1-score) measure the prediction accuracy of every pixel. Unlike most past works that directly compare the prediction results with the ground truth, in this work we use the relaxed version of these metrics following [9], [11], [14].

Let G_{pre} denote the predicted graph and G_{gt} denote the ground-truth graph. Note that both graphs have been densified. Precision is the ratio of pixels in G_{pre} that fall within τ distance to G_{gt} . Recall is the ratio of pixels in G_{gt} that fall within τ distance to G_{pre} . Distance τ reflects the tolerance of inaccuracy. If $\tau = 1$, the relaxed pixel-level metrics degenerate into the commonly-used hard pixel-level metrics. In this paper, we report the evaluation results with τ as 2, 5 and 10 pixels, respectively.

B. The proposed entropy-based connectivity metric (ECM)

Assume the predicted road boundary instances are $G_{pre} = \{G_{pre}^j\}_{j=1}^M$ and the ground-truth instances are $G_{gt} = \{G_{gt}^i\}_{i=1}^N$. In our task, topology errors mainly refer to disconnections. Therefore, the topology correctness metric should suffice the following principles: (1) Punish incorrect disconnections; (2) Assign shorter ground-truth instances G_{gt}^i less weights; (3) Longer incorrect disconnection receives a larger penalty; (4) The prediction whose dominant G_{pre}^j has higher dominance should receive a higher score. Dominant predicted instance G_{pre}^j is the instance with the highest dominance value, where *dominance value* is the ratio of the length of G_{pre}^j to the sum of the length of predicted instances assigned to G_{gt}^i .

Among metrics used in the past work, the widely applied path-based metrics, like APLS and TLTS, meet almost all aforementioned requirements. However, they suffer from randomness and inefficiency due to the random sampling process. Because of this, in [14], an alternative naive connectivity metric is proposed. The authors first match instances in G_{pre} with instances in G_{gt} by minimizing the Hausdorff distance [22]. If an instance




Fig. 5: Comparison of topology metrics. (a) The ground truth; (b)-(f) are predictions with different incorrect disconnections. The evaluation score of each metric is shown in prediction sub-figures. (b)(c) reflects whether metrics suffice principle 1: (c) has more disconnections, thus the connectivity score of (b) should be larger than that of (c). All three metrics meet the requirement. Similarly, (b)(d), (b)(e), (b)(f) reflect whether metrics suffice principle 2,3 and 4, respectively. From the comparison, we find that both APLS and ECM work well for all cases, while the naive connectivity metric and pixel-level metrics fail to follow the principles. Please zoom in for details.

G_{pre}^j matches with G_{gt}^i , G_{pre}^j is assigned to G_{gt}^i . For each ground-truth instance G_{gt}^i , let M_i denote the number of assigned predicted instances. Then the connectivity of instance G_{gt}^i is $C_i = \frac{1(M_i > 0)}{M_i}$, and the final connectivity of the whole patch is the average sum of C_i of all ground-truth instances.

This naive metric can reflect the connectivity of the prediction to some extent, but it fails to meet most aforementioned principles. To relieve this problem, in our previous work *iCurb* [11], we made some adjustments to the naive metric. The matching method is changed to a voting scheme. Each pixel of G_{pre}^j finds its Euclidean-closest G_{gt}^i and makes a vote, then G_{pre}^j is assigned to the ground-truth instance G_{gt}^i that wins the most votes. This voting scheme is more stable than the Hausdorff distance that is sensitive to noises. In addition, a weighting hyper-parameter that gives the longer G_{gt}^i larger weights is assigned to each G_{gt}^i . However, the connectivity metric from the *iCurb* paper cannot meet all the requirements.

Therefore, in this paper, we further modify the connectivity metric and propose an entropy-based connectivity metric (ECM), which is shown in the following equation:

$$ECM = \sum_{i=1}^N \alpha_i e^{-C_i}, \quad \text{where } C_i = \sum_{j=1}^{M_i} -p_j \log(p_j), \quad (1)$$

where C_i is the connectivity of the i -th ground-truth instance G_{gt}^i ; α_i is the completion of G_{gt}^i , which is equal to the sum of the length of assigned instances in G_{pre} projected onto G_{gt}^i divided by the length of G_{gt}^i ; N is the total number of G_{gt}^i in the current patch; M_i is the number of predicted road

boundary instances G_{pre}^j that are assigned to G_{gt}^i , and p_j is the dominance of G_{pre}^j . Entropy is a definition created by Shannon in information theory to measure information content. An event with greater uncertainty tends to have larger entropy. If the dominant predicted instance G_{pre}^{j0} has large dominance value, G_{gt}^i should have good connectivity measure C_i since G_{pre}^{j0} can well approximate G_{gt}^i with more certainty. In summary, ECM assigns longer predicted graph instances with larger weights to calculate C_i , so that short instance could not greatly affect the final evaluation results. Therefore, ECM shows better ability to handle noises and outliers. An example comparing different connectivity metrics is shown in Fig. 5.

V. BASELINE MODELS

In this work, we implement several segmentation-based baselines and graph-based baselines. We also design a new imitation-learning-based method which is enhanced from our previous work *iCurb* [11].

A. Segmentation-based baselines

1) *Naive baseline*: This baseline is proposed by ourselves. Firstly, we employ U-net [23] to obtain the segmentation map of road boundaries. Then, the segmentation results are refined by filtering noisy segments. Finally, the refined segmentation map is skeletonized to get the graph for the road boundaries.

2) *DeepRoadMapper*: This baseline (ICCV2017) [5] proposes to fix disconnections in road network predictions. Due to the similarity between our task and road network detection, we adapt their method for road-boundary detection.

3) *OrientationRefine*: This baseline (CVPR2019) [6] utilizes the orientation map to enhance the segmentation of road networks and achieves improvements. Besides, an extra network is trained to refine the segmentation results iteratively. This work can be directly used for our task without many modifications.

B. Graph-based baselines

1) *RoadTracer*: This baseline (CVPR2018) [7] is believed to be the first work that solves the line-shaped object detection problem by an iterative graph generation approach. However, the network for feature extraction is merely a multi-layer CNN without skip connections.

2) *VecRoad*: This baseline (CVPR2020) [8] greatly improves *RoadTracer*, but follows the same core idea. Res2Net [24] is utilized as the backbone for multi-scale feature extraction. However, there is a limited improvement to the training strategy. Both *VecRoad* and *RoadTracer* work on road network detection, which has different topological characteristics with the road boundary (i.e., road boundaries are scattered polylines without branches while the road network is a connected graph with many intersections), thus we modified their exploration algorithm and the overall pipeline to make them applicable for our task. Besides, road network detection only requires coarse detection of the road centerline, while road boundary detection expects the fine structure of both sides of the road. Therefore, our task has higher demands compared with road network detection.

3) *ConvBoundary*: This baseline (CVPR2019) [14] is the only graph-based work that directly focuses on road-boundary detection. It is a two-stage solution. First, it predicts the inverse-distance map, endpoint map and direction map simultaneously. Then, based on these three maps, a model named cSnake is trained to generate the final graph vertex-by-vertex.

4) *DagMapper*: This baseline (ICCV2019) [9] is designed for road-lane detection. The input of both *DagMapper* and *ConvBoundary* are BEV images of pre-built point-cloud map. Compared with aerial images, BEV images of point-cloud map has much higher resolution and simpler scenarios. For example, the occlusion issue in the point-cloud map is relieved and only the near-road area is covered, which could be regarded as a kind of attention. However, the point-cloud map is time-consuming and expensive to obtain and update, while aerial images are more and more publicly available all over the world. Therefore, our task is more suitable for wide-area applications.

5) *iCurb*: This baseline (RAL2021) [11] is the first work that solves the iterative graph generation from the perspective of imitation learning. *iCurb* has a DAGger-based [25] training strategy, which greatly enhances the performance. This work focuses on road-curb detection, which is a subclass of road boundary, thus it can be directly applied to our task.

Except *RoadTracer*, the code of all the graph-based baselines are not publicly available. So we implement them by ourselves.

C. The Proposed imitation-learning-based baseline

We propose a new imitation-learning-based solution based on our previous work *iCurb* [11], and the new solution is named as *enhanced-iCurb*. To the best of our knowledge, *iCurb* is the first work to solve the line-shaped object detection from the perspective of imitation learning. In [11], a DAGger-based solution is proposed. For each patch, the solution runs a round of restricted exploration as well as N rounds of free exploration, and aggregates the training dataset using the generated samples. Let π^* denote the expert policy and $\hat{\pi}$ denote the learner policy. The task is to learn the $\hat{\pi}$ to mimic π^* . Let \hat{v}_t and v_t^* respectively denote the actions produced by $\hat{\pi}$ and π^* at time t . During the sampling period, v_t is used to denote the vertex to update the graph.

With the obtained prediction \hat{v}_t , we set the closest pixel of the ground-truth road boundary to \hat{v}_t as the label v_t^* to train *iCurb*. This algorithm generates the label on-the-fly and can ensure *iCurb* not be affected by the *teacher-forcing* problem. However, v_t^* heavily relies on \hat{v}_t and does not have a unique value, thus $\hat{\pi}$ may make unpredictable actions and may converge to a sub-optimal policy. In our experiments, for example, we find that the edge length of the road boundary predicted by *iCurb* is almost random sometimes. It requires careful parameter tuning to prevent this. To make the training process more stable and predictable, the orientation map M_O is leveraged to calculate v_t^* and v_t^* has a unique value. We first obtain the radian r_{t-1} of the previous vertex v_{t-1} , then we find the first vertex of the ground-truth road boundary in M_O whose radian value has large enough difference with r_{t-1} , and this vertex is used as v_t^* to train *iCurb*. The new algorithm does not rely on \hat{v}_t so that it guarantees unique v_t^* , which improves the quality of the final graph. So the label v_t^* generated by *enhanced-iCurb* is more reasonable and effective.



Fig. 6: Qualitative demonstrations. Each row shows an example. The first column is the ground truth (cyan lines); column (b) to (d) are segmentation-based baselines (green lines); column (e) to (i) are graph-based baselines and the last column shows the results of the new proposed imitation-learning-based baseline (orange lines for edges and yellow nodes for vertices). For better visualization, the lines are drawn in a thicker width while they are actually one-pixel width. The figure is best viewed in color. Please zoom in for details.

TABLE I: The quantitative comparative results. The best results are highlighted in bold font. For all the metrics, larger values indicate better performance.

Methods	Precision \uparrow			Recall \uparrow			F1-score \uparrow			Naive \uparrow	APLS \uparrow	ECM \uparrow
	2.0	5.0	10.0	2.0	5.0	10.0	2.0	5.0	10.0			
Naive baseline	0.607	0.890	0.928	0.505	0.736	0.768	0.533	0.778	0.811	0.698	0.577	0.550
Deeproadmapper [5]	0.578	0.854	0.898	0.475	0.694	0.725	0.505	0.740	0.775	0.719	0.615	0.595
OrientationRefine [6]	0.620	0.878	0.913	0.602	0.850	0.884	0.605	0.855	0.888	0.797	0.750	0.756
RoadTracer [7]	0.391	0.707	0.791	0.416	0.743	0.821	0.399	0.718	0.798	0.869	0.739	0.824
VecRoad [8]	0.461	0.769	0.854	0.459	0.752	0.830	0.458	0.756	0.837	0.883	0.756	0.846
ConvBoundary [14]	0.510	0.845	0.934	0.455	0.692	0.752	0.465	0.737	0.805	0.958	0.671	0.786
DAGMapper [9]	0.407	0.751	0.868	0.353	0.649	0.747	0.371	0.684	0.787	0.896	0.679	0.758
iCurb [11]	0.550	0.833	0.890	0.538	0.815	0.873	0.542	0.820	0.877	0.910	0.826	0.889
Enhanced-iCurb	0.560	0.839	0.894	0.542	0.811	0.864	0.549	0.821	0.874	0.925	0.822	0.893

VI. EXPERIMENTAL RESULTS AND DISCUSSIONS

A. Experimental setup

We conduct experiments on a PC with an i7-8700K CPU, an NVIDIA GTX1080Ti GPU, and an RTX3090 GPU. The checkpoint with the best performance on the validation set is selected for inference. To measure the efficiency, we record the training time cost as well as the inference time cost of each baseline, and report their average time cost on each patch. For graph-based baselines, the ground-truth initial vertices added with Gaussian noise are used to start the iterative graph generation process. Considering the trade-off between efficiency and effectiveness, the number of rounds of the free exploration for both *iCurb* and *enhanced-iCurb* is set to 1.

B. Evaluation results

The comparative results are shown in Tab. I. Some qualitative demonstrations are shown in Fig. 6. The average processing time

for efficiency evaluation is reported in Tab. II. As segmentation-based baselines directly optimize on pixel values, they tend to present good F1-scores, even the naive baseline. Segmentation-based baselines take relatively less time than graph-based baselines for training since they do not require an iterative process. However, these baselines do not have satisfactory performance on connectivity, because the spatial information of the patch image cannot be fully leveraged. Compared with the *Naive Baseline*, *DeepRoadMapper* can propose candidate connections to correct some disconnection and enhance the connectivity to some extent. But it is still restricted by the accuracy of the semantic segmentation. Once the segmentation network gives poor results, *DeepRoadMapper* cannot effectively improve the final performance. Moreover, generating correction candidates increases the time cost. *OrientationRefine* iteratively refines the obtained segmentation map in pixel-level with an extra network, which can obtain more concise correction results. Thus *OrientationRefine* achieves good connectivity and outperforms all the other methods on F1-Score.

TABLE II: The time consumption of the baselines. We report the average time taken to process a single image.

Methods	Training	Inference
Naive baseline	0.45 s/image	0.33 s/image
Deeproadmapper [5]	1.32 s/image	1.57 s/image
OrientationRefine [6]	1.14 s/image	0.96 s/image
RoadTracer [7]	6.50 s/image	4.92 s/image
VecRoad [8]	19.89 s/image	14.17 s/image
ConvBoundary [14]	17.47 s/image	25.17 s/image
DAGMapper [9]	8.19 s/image	4.92 s/image
iCurb [11]	6.74 s/image	3.11 s/image
Enhanced-iCurb	7.25 s/image	2.38 s/image

Instead of working on pixels, graph-based baselines generate the graph of road boundaries directly. Thus, they perform well on connectivity. However, due to the sequential process for graph generation, this kind of methods exhibit relatively lower pixel-level accuracy when the patch image has complicated scenarios. The training process of *RoadTracer* is fast since it is a small network. But the performance is inferior to others. *VecRoad* follows the similar idea of *RoadTracer* but the network backbone is replaced with a more powerful one, so its performance is improved but the training efficiency is degraded. Even though *ConvBoundary* and *DagMapper* present good results with point-cloud maps on their original tasks, they do not present satisfactory performance on this task (i.e., road-boundary detection using aerial images). *iCurb* provides a solution to this task using imitation learning. It gives better F1-Score and connectivity thanks to the DAGger-based training strategy. *Enhanced-iCurb* utilizes different algorithms to generate the training label and update the graph, which is more stable and predictable. It has good performance on all metrics and qualitative visualizations, thus the superiority of our new proposed method is demonstrated.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a publicly available benchmark dataset named *Topo-boundary* for topological road-boundary detection using BEV aerial images. *Topo-boundary* has 25,295 patches. Each patch consists of a 4-channel aerial image and 8 labels for different deep-learning tasks. We also designed a new metric for better connectivity evaluation. It was employed to compare 9 baselines together with some metrics used in the past works. The dataset and our implemented code for the baselines were publicly available on our project page. In the future, we plan to assign the prediction difficulty scores (i.e., smaller values for easy cases, and larger values for hard cases) to further label the dataset, so that the networks can be trained to be well adapted to various scenarios. Moreover, the patches removed in this paper will be considered for real-world applications.

REFERENCES

[1] X. Lu, Y. Ai, and B. Tian, "Real-time mine road boundary detection and tracking for autonomous truck," *Sensors*, vol. 20, no. 4, p. 1121, 2020.
 [2] X. Zhu, M. Gao, and S. Li, "A real-time road boundary detection algorithm based on driverless cars," in *2015 4th International Conference on Electrical, Electronics and Computer Engineering*. Atlantis Press, 2015, pp. 843–848.

[3] P. Sun, X. Zhao, Z. Xu, R. Wang, and H. Min, "A 3d lidar data-based dedicated road boundary detection algorithm for autonomous vehicles," *IEEE Access*, vol. 7, pp. 29 623–29 638, 2019.
 [4] H. Wang, Y. Sun, and M. Liu, "Self-supervised drivable area and road anomaly segmentation using rgb-d data for robotic wheelchairs," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4386–4393, 2019.
 [5] G. Mátyus, W. Luo, and R. Urtasun, "Deeproadmapper: Extracting road topology from aerial images," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3438–3446.
 [6] A. Batra, S. Singh, G. Pang, S. Basu, C. Jawahar, and M. Paluri, "Improved road connectivity by joint learning of orientation and segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 385–10 393.
 [7] F. Bastani, S. He, S. Abbar, M. Alizadeh, H. Balakrishnan, S. Chawla, S. Madden, and D. DeWitt, "Roadtracer: Automatic extraction of road networks from aerial images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4720–4728.
 [8] Y.-Q. Tan, S.-H. Gao, X.-Y. Li, M.-M. Cheng, and B. Ren, "Vecroad: Point-based iterative graph exploration for road graphs extraction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8910–8918.
 [9] N. Homayounfar, W.-C. Ma, J. Liang, X. Wu, J. Fan, and R. Urtasun, "Dagmapper: Learning to map by discovering lane topology," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 2911–2920.
 [10] D. Belli and T. Kipf, "Image-conditioned graph generation for road network extraction," *NeurIPS 2019 workshop on Graph Representation Learning*, 2019.
 [11] Z. Xu, Y. Sun, and M. Liu, "icurb: Imitation learning-based detection of road curbs using aerial images for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1097–1104, 2021.
 [12] N. O. Department of Information Technology & Telecommunications (DoITT), "NYC-Planimetrics Database," <https://github.com/CityOfNewYork/nyc-planimetrics>, 2019.
 [13] A. Van Etten, D. Lindenbaum, and T. M. Bacastow, "Spacenet: A remote sensing dataset and challenge series," *arXiv preprint arXiv:1807.01232*, 2018.
 [14] J. Liang, N. Homayounfar, W.-C. Ma, S. Wang, and R. Urtasun, "Convolutional recurrent network for road boundary extraction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9512–9521.
 [15] V. Mnih and G. E. Hinton, "Learning to detect roads in high-resolution aerial images," in *European Conference on Computer Vision*. Springer, 2010, pp. 210–223.
 [16] N. Homayounfar, W.-C. Ma, S. Kowshika Lakshmikanth, and R. Urtasun, "Hierarchical recurrent attention networks for structured online maps," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3417–3426.
 [17] Z. Li, J. D. Wegner, and A. Lucchi, "Topological map extraction from overhead images," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1715–1724.
 [18] J. D. Wegner, J. A. Montoya-Zegarra, and K. Schindler, "A higher-order crf model for road network extraction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 1698–1705.
 [19] M. Haklay and P. Weber, "Openstreetmap: User-generated street maps," *IEEE Pervasive computing*, vol. 7, no. 4, pp. 12–18, 2008.
 [20] M. Ranzato, S. Chopra, M. Auli, and W. Zaremba, "Sequence level training with recurrent neural networks," *arXiv preprint arXiv:1511.06732*, 2015.
 [21] J. Liang and R. Urtasun, "End-to-end deep structured models for drawing crosswalks," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 396–412.
 [22] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, "Comparing images using the hausdorff distance," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 15, no. 9, pp. 850–863, 1993.
 [23] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
 [24] S. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. H. Torr, "Res2net: A new multi-scale backbone architecture," *IEEE transactions on pattern analysis and machine intelligence*, 2019.
 [25] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2011, pp. 627–635.