




International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

An AI-Driven Context-Aware Framework For Visibility-Adaptive Road Network Reconstruction and Shortest-Path Healing

Sandhiya S¹, Angela Deepa V R²

¹ Assistant Professor, Department of Computer Science, Women's Christian College, Chennai. Email: sandhiya@wcc.edu.in

² Faculty, Westford Education, India. Email: rians.bliss@gmail.com

Abstract

Satellite images can offer an essential input for road network mapping and applications including navigation, city planning, disaster management, and geographic analysis. Roads can be obstructed from view due to various obstacles like trees, shadows, buildings, clouds, and other objects that cause fragmented roads and disconnected road networks. Road-healing techniques have been using geometrical considerations like distance, direction, and cost of path up until now, but there may be mistakes because reasons for the gaps in the roads are ignored.

In order to solve this problem, this study proposes a context-aware framework utilizing AI for robust reconstruction of the road network as well as repairing of shortest paths in the presence of occlusions. This framework takes into account the context information provided by the images along with graph-theoretic inference for determining whether certain gaps are due to occlusions or they are actually road breaks. An AI or ML-based occlusion context analysis module is included as an extension to the current Context-Aware Reconnection Score (CARS), which allows context features around the gap candidates to be taken in conjunction with geometric features such as endpoint distance, alignment, and road width consistency.

The framework is assessed on the Deep Globe Road Extraction benchmark which consists of 6,226 satellite image tiles and 74,712 annotated gap positions. The use of controlled synthetic occlusions is required to provide ground truth labels for reconnection and quantitative assessment. The best existing CARS guided shortest path configuration managed to obtain an F1 score of 0.448 on the entire test set, beating the existing distance and angle based, geometric shortest path approach and CARS guided configuration alone. The proposed AI extension framework opens up the possibility of learning representations while utilizing the shortest path framework.

Keywords—artificial intelligence, machine learning, satellite imagery, road network reconstruction, occlusion detection, context-aware gap healing, shortest-path optimization, road connectivity.

I. INTRODUCTION

Satellite images are widely used for creating road maps and supporting applications such as urban planning, disaster management, navigation, and geographic analysis. Deep learning-based methods have improved road detection from satellite imagery. However, portions of roads may remain hidden because of trees, building shadows, clouds, vehicles, or other objects. When these regions are not correctly represented during road extraction, the resulting road mask can contain broken or disconnected segments even when the actual road is continuous. This fragmentation is important because a road network used for routing, emergency-response planning, or connectivity analysis must be topologically consistent. A map with numerous gaps may not support reliable shortest-path computation even when pixel-level road segmentation is accurate. Therefore, road-network healing is an important post-processing stage in practical road-extraction pipelines.

Existing healing methods commonly rely on geometric continuity cues, including the distance between fragmented endpoints, directional alignment, road width, and path cost. Although these features provide useful information, geometry alone may not explain why a gap occurs. A short gap caused by tree cover or shadow may require reconnection, whereas a geometrically similar gap caused by a river, building, or other physical boundary should remain disconnected.

The present work builds on the Context-Aware Reconnection Score (CARS), which incorporates contextual information surrounding a candidate road gap into graph-based healing. The proposed AI-driven extension aims to learn contextual characteristics from satellite imagery and integrate them with geometric and graph-based reasoning. Rather than replacing shortest-path reasoning, the AI component is intended to provide additional evidence regarding whether a candidate gap is likely to represent an occluded continuation of a road.

The main contributions are as follows:

- 1) an AI-driven context-aware framework for occlusion-robust road-network reconstruction;
- 2) integration of learned occlusion-context information with geometric road continuity features;
- 3) incorporation of contextual information into shortest-path graph healing;
- 4) a controlled evaluation protocol using synthetically generated occlusion-like gaps with known reconnection labels; and
- 5) evaluation on the complete DeepGlobe road-extraction dataset using precision, recall, F1-score, ablation analysis, and statistical testing.

II. RELATED WORK

A. Road Extraction from Satellite Imagery

Road extraction from satellite imagery is commonly formulated as a semantic segmentation problem in which a learning model identifies pixels belonging to roads. Early neural-network approaches demonstrated the feasibility of automated road detection, while later architectures such as U-Net, SegNet, PSPNet, and DeepLabV3+ improved segmentation by exploiting multiscale and encoder-decoder representations. The DeepGlobe Road Extraction Challenge provided a benchmark for evaluating road extraction methods. D-LinkNet and related architectures further demonstrated the value of pretrained encoders and dilated convolutions for capturing contextual road information.

Although these methods can achieve strong pixel-level segmentation, accurate segmentation does not necessarily produce a complete connected road network. Occlusions and weak visual evidence can result in gaps between road segments, motivating a separate graph-healing stage.

B. Road Connectivity and Graph Healing

Several approaches have been developed to reconnect fragmented road segments. Dynamic-programming approaches can trace centerlines through regions where road evidence is weak. DeepRoadMapper introduced shortest-path reasoning to find possible connections between separated road segments. Other approaches use orientation, segmentation information, connectivity attention, global attention, or graph-based representations to improve road connectivity.

Graph-based methods provide stronger network-level reasoning because road segments can be represented as vertices and edges rather than independent pixels. However, geometric connectivity alone may incorrectly connect segments when an apparent gap corresponds to a genuine physical boundary. This motivates combining graph reasoning with information describing the visual condition surrounding the gap.

C. AI/ML-Based Contextual Reasoning

Machine learning provides a means of learning relationships between image characteristics and reconnection decisions. In the proposed extension, candidate gap regions can be represented using geometric and image-derived contextual features. An AI/ML classifier can learn from controlled positive and negative gap examples to estimate whether a candidate gap is likely to result from occlusion or a genuine road discontinuity. The learned contextual probability can then be integrated into the reconnection cost rather than being used as an isolated decision. This design preserves the interpretable graph-healing component while allowing contextual evidence to be learned from data. The AI component therefore complements, rather than replaces, shortest-path reasoning.

D. Research Gap

The reviewed approaches demonstrate substantial progress in road segmentation, connectivity improvement, and graph construction. However, road gaps caused by occlusion and gaps caused by genuine physical boundaries can have similar geometric characteristics. Existing approaches therefore provide an opportunity for a framework that explicitly learns contextual evidence about the gap and combines it with graph-based reasoning.

A further challenge is the absence of directly observable ground truth for naturally occluded road gaps. The present work addresses this evaluation problem by generating controlled synthetic gaps from known road masks and assigning reconnect or do-not-reconnect labels.

III. PROPOSED METHODOLOGY

A. System Overview

The proposed framework consists of six major stages: satellite-image and road-mask preparation, road skeletonization and graph construction, candidate gap identification, AI-based occlusion-context analysis, context-aware shortest-path healing, and final road-network reconstruction.

The processing flow is:

Satellite Image → Road Segmentation/Mask → Skeletonization → Graph Construction → Candidate Gap Detection → AI/ML Occlusion-Context Analysis → Context-Aware Reconnection → Shortest-Path Healing → Reconstructed Road Network.

B. Problem Formulation

Given a fragmented binary road mask, the mask is skeletonized to obtain a one-pixel-wide centerline representation. A graph $G=(V,E)$ is constructed in which vertices represent junctions and endpoints and edges represent road segments. Let $D \subseteq V$ denote the set of dangling endpoints, defined as vertices connected to only one road segment. Pairs of dangling endpoints within a selected distance limit are treated as candidate locations for missing road connections.

For each candidate pair (a,b) , the framework evaluates geometric continuity and contextual evidence. The final decision is based on the estimated plausibility of a valid road connection and the cost of alternative paths in the existing graph.

C. Geometric Context Features

For each candidate endpoint pair, the framework retains the geometric features used in the original CARS formulation: endpoint distance, angular alignment, and local road-width consistency. These features describe whether two fragmented road segments are spatially and structurally compatible.

The local direction at each endpoint is estimated from the terminal skeleton pixels. Angular deviation is measured between each endpoint direction and the line joining the two endpoints. Distance similarity decreases as the endpoint separation increases, while width consistency measures the similarity between local road widths.

D. AI-Based Occlusion Context Module

The AI extension introduces a learned occlusion-context module for candidate road gaps. For each candidate gap, an image patch or sampled region around the gap is represented using contextual features derived from the satellite imagery. The feature set may include vegetation characteristics, shadow/darkness information, water-like appearance, local texture, edge strength, and the geometric characteristics of the two road endpoints. The training data are generated from known road masks by synthetically removing road segments. Gaps corresponding to simulated tree-canopy or shadow occlusion are assigned positive reconnection labels, whereas gaps representing genuine physical boundaries such as rivers or buildings are assigned negative labels. This controlled construction provides a supervised learning problem with known ground truth.

The AI/ML model produces an occlusion-context probability $P_{occ}(a,b)$, representing the likelihood that the candidate gap is an occlusion-induced missing connection. This learned probability can replace or complement the hand-designed contextual statistics used in the original CARS formulation.

E. Context-Aware Reconnection Score

The geometric similarity components and the learned contextual probability are combined to produce a context-aware reconnection score. A generic formulation is:

$$CARS_AI(a,b) = F(sdist, sangle, swidth, P_{occ})$$

where $sdist$ denotes distance similarity, $sangle$ denotes angular similarity, $swidth$ denotes road-width consistency, and P_{occ} denotes the learned probability of occlusion.

The exact model and weights should be selected using only the development subset and then fixed before final evaluation. This prevents test-set information from influencing model selection.

F. Shortest-Path Healing

A geometric shortest-path healing method is used as a strong graph-based comparison. Candidate reconnections are represented as weighted virtual edges. A candidate connection is accepted only when the direct virtual edge satisfies the predefined cost condition and is supported by shortest-path reasoning after temporarily removing the candidate edge.

In the proposed AI-driven framework, the learned occlusion-context information modifies the candidate edge cost. When a gap has a high probability of being caused by occlusion, the contextual penalty is reduced; when the probability is low, the candidate connection receives a higher penalty. Geometric plausibility remains part of the cost so that image-based context does not override road-network structure.

G. Algorithmic Workflow

Step 1: Obtain the satellite image and corresponding road mask.

Step 2: Skeletonize the road mask.

Step 3: Construct the road graph from junctions and endpoints.

Step 4: Identify dangling endpoints and generate candidate endpoint pairs.

Step 5: Extract geometric and image-context features for each candidate gap.

Step 6: Apply the AI/ML occlusion-context model to estimate P_{occ} .

Step 7: Combine learned contextual evidence with geometric similarity.

Step 8: Assign a context-aware cost to each candidate reconnection.


Step 9: Apply shortest-path reasoning to determine plausible connections.

Step 10: Add accepted connections to the graph.

Step 11: Produce the healed road network.

Step 12: Evaluate reconnection precision, recall, and F1-score against the known synthetic ground truth.

Road Network Healing Workflow



IV. EXPERIMENTAL SETUP

A. Dataset

The DeepGlobe Road Extraction dataset is used for evaluation. The dataset contains 6,226 satellite-image tiles with ground-truth binary road masks. Because naturally occurring occlusions do not provide an independently observable correct healed topology, controlled synthetic gaps are generated from the ground-truth road masks.

Each tile contributes eight positive gap sites and four negative gap sites. Positive sites simulate occlusion conditions in which the road should be reconnected, while negative sites simulate genuine physical breaks in which reconnection should not occur. The complete dataset therefore provides 74,712 labeled gap sites for evaluation.

B. Experimental Configurations

The evaluation should compare the following configurations:

- 1) distance-and-angle baseline;
- 2) geometric shortest-path healing;
- 3) original CARS standalone;
- 4) AI/ML-based occlusion-context model;
- 5) AI-enhanced CARS-guided shortest-path healing; and
- 6) ablation variants without learned contextual information.

The AI-specific configurations require training on a development subset and evaluation on held-out data. Parameters selected during development must remain fixed for the final evaluation.

C. Evaluation Metrics

A healing decision is treated as a true positive when it reconnects within 25 pixels of a should-reconnect ground-truth site. A false positive occurs when an incorrect connection is made near a should-not-reconnect site or without a matching ground-truth site. A false negative occurs when a should-reconnect site receives no matching connection.

Precision, recall, and F1-score are used to evaluate reconstruction performance. F1-score provides a combined measure of precision and recall and is particularly useful because incorrect reconnections and missed reconnections both affect road-network quality.

V. RESULTS

The existing implementation evaluated five configurations on the complete DeepGlobe dataset. The reported aggregate results are shown below. These values belong to the existing CARS-based implementation and should not be presented as results of the new AI/ML module until the AI model has been implemented and experimentally evaluated.

Method	Precision	Recall	F1-score
Distance + angle baseline	0.425	0.378	0.400
Shortest-path (geometric only)	0.321	0.664	0.433
CARS standalone, no occlusion	0.243	0.697	0.361
CARS standalone (full)	0.319	0.623	0.422
CARS-guided shortest-path	0.395	0.519	0.448

The existing CARS-guided shortest-path configuration achieved an F1-score of 0.448, compared with 0.433 for the geometric shortest-path method, 0.422 for standalone CARS, and 0.400 for the distance-and-angle baseline. The existing study therefore provides the baseline evidence on which the AI-driven extension can be evaluated.

A. Ablation and Statistical Analysis

The original evaluation reported that removing the occlusion-context component reduced the standalone CARS F1-score from 0.422 to 0.361. Paired statistical tests across the 6,226 tiles also indicated statistically significant differences between selected configurations. The improvement of the final CARS-guided shortest-path method over the plain shortest-path baseline was statistically significant, although the absolute F1 improvement was small. These results support retaining the contextual component while motivating further investigation of learned contextual representations.

B. Required AI/ML Evaluation

To substantiate the AI-driven title, the learned occlusion-context model must be evaluated separately and then as part of the complete healing framework. The following additional results should be reported after implementation: classifier accuracy, precision, recall, F1-score, confusion matrix, and the final road-healing precision, recall, and F1-score. A comparison between the hand-designed CARS contextual features and the learned AI representation should also be included.

The final AI model should be evaluated on held-out tiles that were not used for feature selection, model training, or threshold tuning. This separation is important for demonstrating generalization and avoiding optimistic performance estimates.

VI. DISCUSSION

The existing results indicate that contextual information is more useful when integrated into shortest-path reasoning than when used as an independent pairwise reconnection decision. The geometric shortest-path method achieved an F1-score of 0.433, while standalone CARS achieved 0.422. When the contextual signal was integrated into the shortest-path cost, the resulting configuration achieved 0.448.

The proposed AI extension is motivated by the observation that the original contextual terms are based on hand-designed image statistics. A learned model may capture more complex relationships among vegetation, shadow, water-like appearance, local texture, road geometry, and other visual characteristics. Such learning could improve the estimation of whether a candidate gap is caused by occlusion.

However, the AI component should be evaluated independently rather than assumed to improve performance. If the learned model does not outperform the existing contextual formulation, the result should be reported transparently. The principal contribution would then be the analysis of learned versus hand-designed context within a graph-healing framework.

VII. LIMITATIONS

Several limitations of the current study remain. First, the evaluation uses synthetically introduced gaps rather than naturally occluded imagery because natural occlusions do not provide directly observable ground-truth topology. Second, the existing CARS contextual terms use hand-designed statistics. Third, the reported experiments use ground-truth segmentation masks with synthetic gaps, thereby isolating the healing stage from upstream segmentation errors. Fourth, evaluation is based on a single dataset, and generalization to other satellite-image sources has not yet been established. Finally, the AI extension requires additional experiments to determine whether learned contextual representations provide a measurable improvement over the existing handcrafted formulation.

VIII. CONCLUSION AND FUTURE WORK

This paper presents an AI-driven context-aware framework for occlusion-robust road-network reconstruction and shortest-path healing. The framework combines road-network geometry, image-derived contextual evidence, AI/ML-based occlusion analysis, and graph-based shortest-path reasoning. The existing CARS-guided shortest-path implementation achieved an F1-score of 0.448 on the complete DeepGlobe evaluation set, demonstrating the value of integrating contextual information with graph reasoning.

The next stage of the work is to implement and evaluate the AI/ML occlusion-context module using the controlled positive and negative gap samples. Future research can investigate deep feature extraction, transformer-based image representations, graph neural networks, multi-dataset validation, and naturally occluded satellite imagery. Additional work can also evaluate the effect of segmentation uncertainty and domain variation on road-network reconstruction.

The overall objective is to develop a robust intelligent road-reconstruction system that can use both visual context and network-level reasoning to recover plausible road connectivity while avoiding incorrect connections across genuine physical boundaries.

References

- [1] Mnih and Hinton, neural-network-based road detection from aerial imagery.
- [2] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," 2015.
- [3] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation," 2017.
- [4] H. Zhao et al., "Pyramid Scene Parsing Network," 2017.
- [5] L.-C. Chen et al., "Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation," 2018.

- [6] I. Demir et al., "DeepGlobe 2018: A Challenge to Parse the Earth through Satellite Images," 2018.
- [7] L. Zhou et al., "D-LinkNet: LinkNet with Pretrained Encoder and Dilated Convolution for High Resolution Satellite Imagery Road Extraction," 2018.
- [8] A. Gruen and H. Li, dynamic-programming approaches for road-centerline extraction.
- [9] A. Mena, review of road extraction methods and occlusion challenges.
- [10] M. Mátyus et al., "DeepRoadMapper: Extracting Road Topology from Aerial Images," 2017.
- [11] G. Batra et al., learned connectivity prediction using road orientation and image features.
- [12] C. Mei et al., "CoANet: Connectivity Attention Network for Road Extraction," 2021.
- [13] X. Lu et al., GAMSNet for road detection and connectivity.
- [14] F. Bastani et al., "RoadTracer: Automatic Extraction of Road Networks from Aerial Images," 2018.
- [15] H. He et al., "Sat2Graph: Road Graph Extraction from Aerial Images," 2020.
- [16] Z. Xu et al., RNGDet for road-network vertex detection and construction.
- [17] A. Dosovitskiy et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," 2021.
- [18] L. Freeman, work on betweenness centrality and graph-network analysis.