Benchmarking Road Network Extraction Methods for Power Distribution Spatial Planning

Janez Križaj, Vitomir Štruc
University of Ljubljana, Faculty of Electrical Engineering, Slovenia
janez.krizaj@fe.uni-lj.si, vitomir.struc@fe.uni-lj.si

Accurate road network extraction from high-resolution overhead imagery is a prerequisite for corridor-aware routing and cost modelling in power-system infrastructure planning, where medium- and low-voltage lines often follow transport rights-of-way. Yet roads are thin, occluded, and cluttered, and downstream optimisation needs routable graphs, not just masks. We benchmark six representative pipelines (SAM-Road, D-LinkNet, CRESI, CU-dGCN, Sat2Graph, and U-Net-ResNet18) on SpaceNet imagery. The comparison reveals trade-offs between segmentation-centric and graph-aware designs and highlights where topology fails most often. Finally, practical guidance is provided for selecting road extractors that deliver reliable, optimisation-ready networks for power distribution corridor planning and the routing of medium- and low-voltage lines.

DOI
https://doi.org/10.18690/um.feri.4.2026.3

ISBN
978-961-299-111-1

Keywords:
road network extraction, road graph extraction, satellite imagery, benchmarking, network connectivity, power distribution spatial planning

University of Maribor Press

DOI
https://doi.org/
10.18690/um.feri.4.2026.3

ISBN
978-961-299-111-1

Ključne besede:
ekstrakcija cestnega omrežja,
satelitske slike,
primerjalno ovrednotenje,
analiza povezljivosti,
prostorsko načrtovanje distribucijskega elektroenergetskega omrežja

Univerzitetna založba
Univerze v Mariboru

PRIMERJAVA METOD EKSTRAKCIJE CESTNEGA OMREŽJA ZA PROSTORSKO NAČRTOVANJE DISTRIBUCIJE ELEKTRIČNE ENERGIJE

JANEZ KRIŽAJ, VITOMIR ŠTRUC
Univerza v Ljubljani, Fakulteta za elektrotehniko, Slovenija
janez.krizaj@fe.uni-lj.si, vitomir.struc@fe.uni-lj.si

Natančna ekstrakcija cestnega omrežja iz satelitskih posnetkov omogoča določanje poteka in modeliranje stroškov pri načrtovanju infrastrukture elektroenergetskega sistema, kjer srednje- in nizkonapetostni koridorji pogosto sledijo cestam. Ker so ceste na satelitskih slikah pogosto ozke, deloma prikrite in nejasne, lahko manjše nepravilnosti v segmentacijskih maskah po vektorski pretvorbi povzročijo prekinjeno povezljivost, zato je za načrtovanje bolj koristno uporabiti grafe, ki omogočajo načrtovanje poti. Na zbirki SpaceNet smo primerjali šest reprezentativnih postopkov (SAM-Road, D-LinkNet, CRESI, CU-dGCN, Sat2Graph in U-Net--ResNet18). V primerjavi proučimo razlike med klasičnimi segmentacijskimi postopki in novejšimi postopki z grafnimi zasnovami ter analiziramo, kje posamezni postopki najpogosteje odpovejo. Navedemo tudi praktična navodila za izbiro ustreznih postopkov, ki zagotavljajo zanesljivo detekcijo cest, primerno za optimizacijo pri načrtovanju srednje- in nizko-napetostninih vodov v distribucijskem električnem omrežju.

1 Introduction

Accurate road network extraction from aerial and satellite imagery is a longstanding challenge in computer vision and remote sensing. It underpins crucial applications in navigation, autonomous driving, urban planning, and disaster response, and it is also a key input for the spatial planning of power system infrastructure. In distribution network development, candidate medium- and low-voltage corridors are often derived from transport right-of-way: routing along roads can reduce permitting effort, limit excavation on private parcels, and guarantee construction and maintenance access. Consequently, the quality and completeness of the road layer can directly affect corridor-aware routing, cost modelling, and feasibility screening of candidate alignments.

While early approaches relied on manual mapping or classical image processing, the advent of deep learning has led to significant progress in automated road extraction. Yet, producing routable road networks – continuous centerline graphs suitable for path planning – remains difficult. CNN-based segmentation often yields fragmented or noisy road masks, necessitating post-processing (e.g. thinning or connectivity analysis) to derive vector road graphs. Importantly, errors that are visually minor at the pixel level (small gaps, slightly misplaced intersections) can become major issues once converted into graphs, by disconnecting otherwise feasible corridors or introducing spurious shortcuts.

Motivated by these challenges, this paper conducts a comparative evaluation of modern road extraction techniques, covering both pixel-level and topology-aware paradigms. We focus on state-of-the-art models that span the spectrum from pure mask prediction to approaches that explicitly model connectivity, with the goal of clarifying how structural priors affect accuracy and practical suitability for optimisation-ready workflows in infrastructure planning.

The contributions of this work are as follows:

– Clear taxonomy of road extraction methods: We outline a taxonomy dividing approaches into segmentation-based, skeleton-based, and graph-based methods, informed by recent surveys (Chen et al., 2022) (Sharma et al., 2024) (Liu et al., 2024) (Wang et al., 2025). This categorization highlights the field's evolution from pixel-level segmentation towards representations that capture road connectivity and topology.

- Systematic evaluation of representative models: We implement and evaluate six representative pipelines: D-LinkNet (Zhou et al., 2018) and U-Net–ResNet18 (Robinson et al., 2024) as pure segmentation baselines, CU-dGCN (Vekinis, 2023) as a segmentation model augmented with graph reasoning, CRESI (Etten, 2020) as a segmentation+skeleton pipeline, Sat2Graph (He et al., 2020) as a direct graph prediction approach, and SAM-Road (Hetang et al., 2024) as a foundation-model-based method that combines mask prediction with a lightweight graph module.
- Insights into topology-aware extraction: Through qualitative and quantitative analysis, we examine the ability of each method to produce coherent road networks. We show how methods that leverage graph structure (either via architectural components or post-processing) improve topological correctness (fewer broken segments and spuriously disconnected roads) compared to vanilla segmentation.
- Practical considerations for spatial optimisation: We relate the reported road-network metrics to requirements that arise in corridor-aware planning (e.g., avoiding disconnected corridors and unrealistic shortcuts), with emphasis on the needs of power system infrastructure planning.

Overall, our work demonstrates how increasingly sophisticated models address the limitations of pure segmentation, and it offers guidance for selecting road extraction pipelines when the end goal is a connected, routable network for downstream planning workflows. The next sections present related work, describe the evaluated methods, and report experimental results.

2 Related Work

Road extraction from remote sensing imagery has been widely studied, and recent surveys document a clear shift from classical image processing toward deep learning pipelines (Chen et al., 2022) (Sharma et al., 2024) (Liu et al., 2024) (Wang et al., 2025). Consistent with these surveys, we adopt a taxonomy that groups approaches into: (*i*) segmentation-based methods that predict pixel-wise road masks, (*ii*) skeleton-based pipelines that derive centerlines/graphs from predicted masks via thinning and pruning, and (*iii*) graph-based approaches that model connectivity explicitly, either by learning on graphs or by directly predicting graph structures. This taxonomy is practically relevant whenever the downstream goal is a routable network: small gaps or spurious links that are visually minor at the pixel level can disconnect feasible routes or introduce unrealistic shortcuts once a mask is converted into a graph.

Segmentation-based approaches commonly cast road extraction as semantic segmentation with encoder--decoder CNNs (e.g., FCN/U-Net/SegNet families) (Badrinarayanan et al., 2017) (Chen et al., 2018), producing a dense road mask that can be thresholded and later vectorized. D-LinkNet (Zhou et al., 2018) improves context aggregation (e.g., via dilated convolutions), while lightweight variants such as U-Net--ResNet18 (Robinson et al., 2024) remain strong baselines. A key limitation for optimisation-ready use is that visually minor mask defects (gaps, junction artifacts) can become major topological errors after skeletonization.

Skeleton-based pipelines address this limitation by augmenting segmentation with explicit mask cleanup, skeletonization, and pruning to improve routability; a representative example is CRESI (Etten, 2020). While effective in practice, their final graph quality remains sensitive to the continuity of the predicted mask and to heuristic post-processing choices.

Graph-based methods target connectivity more directly. Some incorporate topology cues into learning, for example via multi-task designs that add structural supervision such as orientations (Batra et al., 2019), while hybrids inject graph reasoning modules to propagate connectivity information (Vekinis, 2023). In parallel, graph-centric approaches avoid heuristic vectorization by predicting road graphs explicitly: RoadTracer (Bastani et al., 2018) pioneered iterative graph construction, Sat2Graph (He et al., 2020) predicts vertices and edges in a single forward pass, and transformer-based detectors such as RNGDet (Xu et al., 2022) and RNGDet++ (Xu et al., 2023) extend this paradigm with detection-style graph decoding. This trend toward graph outputs (rather than raster masks) is also emphasized in recent surveys (Chen et al., 2022) (Liu et al., 2024) (Wang et al., 2025).

3 Methods

We compare six representative approaches spanning the three categories introduced above: segmentation-based mask predictors (D-LinkNet, UNet-ResNet18), a skeleton-based pipeline that derives a centerline graph from a predicted mask

(CRESI), and graph-based methods that explicitly model connectivity - either by refining mask predictions with learned graph reasoning (CU-dGCN, SAM-Road) or by directly predicting a road graph (Sat2Graph). Below we summarize each model's architecture and output format.

3.1 U-Net–ResNet18 (Segmentation baseline)

U-Net–ResNet18 (Robinson et al., 2024) employs a standard U-Net encoder-decoder architecture with a ResNet-18 backbone, providing a lightweight yet effective road segmentation baseline. As a pure segmentation approach, it outputs a binary road mask after thresholding; connectivity is only implicitly learned and not explicitly enforced.

3.2 D-LinkNet (Segmentation baseline)

D-LinkNet (Zhou et al., 2018) is a CNN road segmenter based on an encoder-decoder architecture with skip connections. Dilated convolutions enlarge the receptive field while preserving spatial detail, which helps bridge gaps caused by occlusions or low contrast. In typical configurations, a pretrained ResNet backbone extracts multi-scale features and the decoder fuses them into a dense road-probability map. The output is a binary mask after thresholding; connectivity is encouraged only implicitly through context aggregation and is not explicitly enforced.

3.3 CRESI (Segmentation + skeleton)

CRESI (Etten, 2020) targets large-scale road graph extraction through a pragmatic multi-stage pipeline. A segmentation model predicts road likelihood on overlapping tiles, which are then mosaicked into a global mask to reduce boundary artifacts. Post-processing cleans the mask and fills small gaps before skeletonization. The skeleton is converted into a graph, and pruning removes spurs and short branches to improve routability. This design emphasizes that the graph-construction stage can be as important as the segmentation backbone.

3.4 CU-dGCN (Segmentation + graph reasoning)

CU-dGCN (Vekinis, 2023) couples convolutional image features with graph neural network reasoning. A road centerline graph is constructed or predicted and then refined via message passing, allowing connectivity cues to propagate across gaps and complex junctions. The refined graph representation is mapped back to improve the road prediction, effectively combining local appearance with global structure. This hybrid design aims to deliver better routability without fully abandoning the efficiency and stability of mask-based segmentation.

3.5 SAM-Road (Segmentation + graph reasoning)

SAM-Road (Hetang et al., 2024) adapts the Segment Anything Model (SAM) to road extraction using road-specific prompting strategies and a decoding pipeline that yields road masks and graph representations. Leveraging a foundation-model prior can help under domain shift and with limited labeled data. In practice, SAM-Road still depends on prompt design and decoding choices to avoid fragmented masks and to produce a routable network, but it reduces reliance on training a task-specific segmenter from scratch.

3.6 Sat2Graph (Direct graph prediction)

Sat2Graph (He et al., 2020) predicts a road graph in a single forward pass by learning structured representations of vertices and edges from the image. Compared to tracing, global prediction is less sensitive to local dead-ends and can better preserve long-range continuity. The approach relies on carefully defined training targets and a robust graph decoding stage to avoid missing thin roads or introducing shortcuts. The output is a graph, which can be rasterized to a mask for overlap-based evaluation when needed.

4 Experiments

We compare representative segmentation-first and graph-centric road extraction methods on the SpaceNet roads benchmark, which provides satellite imagery paired with ground-truth vector road graphs. Because our target downstream use is corridor-aware routing for spatial optimisation (including power system infrastructure planning), we focus on connectivity and routability rather than pixel overlap alone. To make the comparison consistent across output types, mask predictions are thresholded and converted into centerline graphs using a standard thinning-and-tracing pipeline with light cleanup, while graph predictions are evaluated directly after minor snapping of near-coincident vertices. We report TOPO[1] precision/recall/F1 and APLS (Average Path Length Similarity) following the SpaceNet evaluation definitions; in Table 1 the precision/recall/F1 values correspond to TOPO rather than pixel-level segmentation scores.

Table 1: SpaceNet results for representative methods. Prec./Rec./F1 are TOPO metrics; APLS measures routing fidelity. Values are in %.

Method	Prec. ↑	Rec. ↑	F1 ↑	APLS ↑
U-Net–ResNet18	68.96	66.32	67.61	53.77
D-LinkNet	82.42	60.06	68.80	56.96
CRESI	79.11	69.47	71.87	56.38
CU-dGCN	62.13	61.66	62.32	58.19
Sat2Graph	85.93	**76.55**	**80.97**	64.43
SAM-Road	**93.03**	70.97	80.52	**71.64**

Table 1 summarizes quantitative results. For corridor selection and cost modelling, APLS is particularly informative because it measures shortest-path similarity and is therefore sensitive to disconnections that would force unrealistic detours in an optimised route. TOPO complements APLS by capturing local structural agreement around intersections and branches: high precision discourages spurious links (which can create infeasible shortcuts along non-existent roads), while high recall reduces the risk of missing feasible corridors. Among the compared methods, SAM-Road achieves the highest APLS and very high TOPO precision, indicating clean graph predictions with comparatively few false connections, but its lower recall suggests that it can miss parts of the network. Sat2Graph provides the best balance in TOPO and also strong APLS, which can be advantageous when the planning goal is to enumerate alternative access corridors rather than only the most certain ones. CU-dGCN and CRESI sit in the mid-range, while the segmentation-first baselines (D-LinkNet and U-Net–ResNet18) exhibit lower APLS, consistent with the fact that small mask gaps and junction errors tend to be amplified during vectorization. Overall, the table suggests that graph-centric approaches are better suited for automation-heavy planning workflows, whereas pure segmentation outputs are

[1] TOPO evaluates the similarity of reachable subgraphs from randomly sampled vertices in the predicted road graph and ground truth.

more likely to require additional post-processing or manual GIS corrections before they can reliably support corridor-aware routing.

Fig. 1 provides a qualitative comparison on a challenging orthophoto from Surveying and Mapping Authority of the republic of Slovenia (hereinafter GURS), where we use the OSM road layer (OpenStreetMap contributors, 2017) as a pseudo ground truth reference for visual assessment. Under this interpretation, segments present in OSM but missing in the prediction correspond to false negatives (potentially removing feasible corridors), whereas predicted segments absent in OSM correspond to false positives (potentially creating spurious, non-existent shortcuts). We note that OSM itself can be incomplete or slightly misaligned, so some apparent false positives may in reality be unmapped roads; nevertheless, for planning workflows that start from an OSM-like baseline, such deviations would still require additional verification.

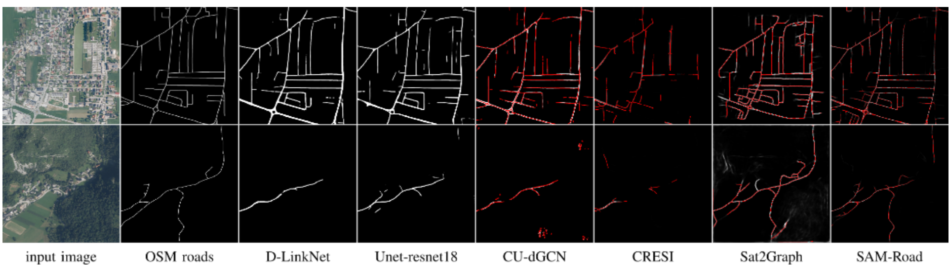

Figure 1: Visual comparison of road extraction outputs. Column 1: GURS orthophoto (GURS, 2026). Column 2: OSM road layer rasterized as a binary mask. Remaining columns: method-predicted road masks; where a road graph/centerline is natively available, it is drawn as red lines over the mask.
Source: GURS and author-generated images.

In the urban part (top of Fig. 1), D-LinkNet largely follows the main OSM corridors but produces widened road regions and small spurious blobs that would likely turn into extra branches after thinning. The U-Net baseline is noisier, deviating from the OSM topology with more fragmented and spurious segments. CU-dGCN and CRESI suppress many of these artifacts and better preserve the mapped connectivity, but they still exhibit breaks where OSM indicates continuous roads, especially near shadows and vegetation boundaries.

In the rural/forested part (bottom of Fig. 1), the OSM road is partially obscured by canopy and overgrowth in the orthophoto. Here, Sat2Graph appears to recover a more continuous connection along the OSM corridor than the other methods, suggesting improved robustness to partial occlusion in the tested GURS imagery. A plausible explanation is that Sat2Graph's direct graph supervision and global decoding impose a stronger connectivity prior and may generalize better under dataset shift between its training data and national orthophotos. SAM-Road yields a clean, high-precision network that aligns well with the main OSM roads, but it can still break under heavy occlusion and may miss minor branches, which is consistent with the precision--recall trade-off observed in Table 1.

5 Conclusion

This paper organized road extraction methods for satellite imagery by output representation and compared representative approaches using graph-oriented evaluation. The main takeaway is that road masks that look plausible are not necessarily routable once converted to a centerline graph: small gaps and junction errors can dominate final network quality and are better captured by TOPO and APLS than by pixel overlap alone. This distinction matters directly for power system spatial planning, where road networks are commonly used as candidate corridors and access constraints for routing underground cables or overhead lines. In this setting, disconnected graphs can eliminate feasible corridors, while spurious links can lead to unrealistic low-cost routes that would later fail during permitting or field verification.

Across the compared methods, the quantitative results indicate that graph-centric approaches (SAM-Road and Sat2Graph) provide higher routing fidelity (APLS) and more stable connectivity than segmentation-first baselines, making them more suitable for automation-heavy corridor-aware optimisation. The qualitative comparison on GURS orthophotos further highlights that OSM roads provide a valuable prior but can be incomplete in rural or rapidly changing areas, and vegetation occlusion remains a key failure mode. In our tested examples, Sat2Graph appears more robust in recovering roads that are partially hidden by overgrowth, which may stem from its direct graph supervision and stronger generalization under domain shift.

Acknowledgement

The authors gratefully acknowledge financial support from the fundamental research project Opti-AI (Optimal Integration of Power System Infrastructure through AI-based Spatial Planning, project no. J2-60030), funded by the Slovenian Research and Innovation Agency (ARIS).

Literatura

Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell., 39*, 2481–2495.

Bastani, F., He, S., Abbar, S., Alizadeh, M., Balakrishnan, H., Chawla, S., Madden, S., & DeWitt, D. (2018). Roadtracer: Automatic extraction of road networks from aerial images. *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, (str. 4720–4728).

Batra, A., Singh, S., Pang, G., Basu, S., Jawahar, C. V., & Paluri, M. (2019). Improved road connectivity by joint learning of orientation and segmentation. *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, (str. 10385–10393).

Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., & Adam, H. (2018). Encoder-decoder with atrous separable convolution for semantic image segmentation. *Proc. Eur. Conf. Comput. Vis. (ECCV)*, (str. 801–818).

Chen, Z., Deng, L., Luo, Y., Li, D., Junior, J. M., Gonçalves, W. N., Nurunnabi, A. A., Li, J., Wang, C., & Li, D. (2022). Road extraction in remote sensing data: A survey. *Int. J. Appl. Earth Obs. Geoinf., 112*, 102833. https://doi.org/https://doi.org/10.1016/j.jag.2022.102833

Etten, A. V. (2020). City-Scale Road Extraction from Satellite Imagery v2: Road Speeds and Travel Times. *2020 IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, (str. 1775-1784). https://doi.org/10.1109/WACV45572.2020.9093593

GURS. (2026). National orthophoto dataset. *Surveying and Mapping Authority of the Republic of Slovenia.* Ministry of Natural Resources and Spatial Planning. https://podatki.gov.si/dataset/ortofoto

He, S., Bastani, F., Jagwani, S., Alizadeh, M., Balakrishnan, H., Chawla, S., Elshrif, M. M., Madden, S., & Sadeghi, M. A. (2020). Sat2Graph: Road Graph Extraction through Graph-Tensor Encoding. *Proc. Eur. Conf. Comput. Vis. (ECCV).* https://api.semanticscholar.org/CorpusID:220647459

Hetang, C., Xue, H., Le, C., Yue, T., Wang, W., & He, Y. (2024). Segment anything model for road network graph extraction. *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, (str. 2556–2566).

Liu, R., Wu, J., Lu, W., Miao, Q., Zhang, H., Liu, X., Lu, Z., & Li, L. (2024). A Review of Deep Learning-Based Methods for Road Extraction from High-Resolution Remote Sensing Images. *Remote Sens., 16*. https://doi.org/10.3390/rs16122056

OpenStreetMap contributors. (2017). Planet dump retrieved from https://planet.osm.org . *Planet dump retrieved from https://planet.osm.org .*

Robinson, C., Corley, I., Ortiz, A., Dodhia, R., Ferres, J. M., & Najafirad, P. (2024). Seeing the roads through the trees: A benchmark for modeling spatial dependencies with aerial imagery. *Seeing the roads through the trees: A benchmark for modeling spatial dependencies with aerial imagery.* https://arxiv.org/abs/2401.06762

Sharma, P., Kumar, R., Gupta, M., & Nayyar, A. (2024). A critical analysis of road network extraction using remote sensing images with deep learning. *Spat. Inf. Res., 32*, 485–495. https://doi.org/10.1007/s41324-024-00576-y

Vekinis, A. A. (2023). Graph Reasoned Multi-Scale Road Segmentation in Remote Sensing Imagery. *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, (str. 6890-6893). https://doi.org/10.1109/IGARSS52108.2023.10281660

Wang, X., Jin, X., Dai, Z., Wu, Y., & Chehri, A. (2025). Deep Learning-Based Methods for Road Extraction From Remote Sensing Images: A vision, survey, and future directions. *IEEE Geosci. Remote Sens. Mag., 13*, 55-78. https://doi.org/10.1109/MGRS.2024.3491014

Xu, Z., Liu, Y., Gan, L., Sun, Y., Wu, X., Liu, M., & Wang, L. (2022). Rngdet: Road network graph detection by transformer in aerial images. *IEEE Trans. Geosci. Remote Sens., 60*, 1–12.

Xu, Z., Liu, Y., Sun, Y., Liu, M., & Wang, L. (2023). Rngdet++: Road network graph detection by transformer with instance segmentation and multi-scale features enhancement. *IEEE Robot. Autom. Lett., 8*, 2991–2998.

Zhou, L., Zhang, C., & Wu, M. (2018). D-LinkNet: LinkNet with Pretrained Encoder and Dilated Convolution for High Resolution Satellite Imagery Road Extraction. *2018 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, (str. 192-1924). https://doi.org/10.1109/CVPRW.2018.00034